PRODUCT UPDATES IN THE PRESENCE OF CONSUMER PROGRESSION

Paulo Albuquerque* Yulia Nevskaya†

Preliminary. Please do not cite without authors' permission

Abstract

We propose a dynamic model to study the managerial decisions regarding the speed of product innovation in categories where there is an increasing involvement of consumers with the product. In our framework, we consider products where the utility of future consumption depends directly on the decision to consume content today. Our approach is applicable to a wide variety of experiential products, such as video games, TV series, and other entertainment products. The demand for product content is modeled as a function of consumer preferences for current content, previous experience with the product, and forward-looking expectations about future content and enjoyment. We use data on consumer participation from the popular online computer game "World of Warcraft" to empirically apply our model. We show that decision to introduce product updates strongly depends on consumer heterogeneity, rate of content consumption and satiation, and social interaction of consumers. We are able to quantify the value of a product update in terms of consumer participation and we provide managerial recommendations to improve scheduling of product innovation.

Keywords: dynamic demand models, forward-looking consumers, online content.

* Assistant Professor of Marketing at the Simon Graduate School of Business, University of Rochester.
† Ph.D. Candidate at the Simon Graduate School of Business, University of Rochester.

1 Introduction

The development and introduction of product updates, as well as their timing, are among the most important decisions of marketing managers (Norton and Bass, 1987). During its lifetime cycle, a product can go through multiple iterations with new and improved versions of the product being introduced to keep consumers interested and to increase repeat purchase and product replacement and usage. More specifically, in categories involving experiential products where consumers frequently use a product or enjoy related content over time, product updates are essential. In such categories, the consumer focuses not only on tangible product characteristics, but also on intangibles related to product usage, such as the sociability of actions and the satisfaction from experiencing previous unexplored content (Holbrook and Hirschman, 1982).

When a product and its content are enjoyed repeatedly by consumers and driven by a hedonic response, consumer loyalty and interest depend on both the benefits of the current experiences and the availability of unexplored content that provides new experiences. Examples of experiential products studied in previous consumer behavior literature include movies, television series, music, video games, vacations, jewelry, etc.. If innovative content is not introduced, a decline in demand is likely to happen due to satiation with old content and the emergence of substitute products. Hence, firms need to provide a constant stream of new material especially when revenues depend on the continuity of content consumption. An anecdotal example of how important product updates are to managers and consumers alike in one of these experiential products was provided in May 2011 at the earnings call of Activision Blizzard, one of the major developers of computer games software. At the call, the popular online game "World of Warcraft" was the major discussion topic since it generates a large percentage of the firm's profits. Subscriptions for the game had declined from 12 million paying customers at the end of 2010 to 11.4 million at the end of March of 2011 and continued to rapidly drop until May. In response to questions about this decline, the company's CEO, Mike Morhaime, said that "subscriber base does not change linearly. It fluctuates based on content consumption, which players seem to be doing a whole lot of – at a more rapid pace", and continued by promising "faster release of new content" to respond to the demand decline.

Experiential products are frequently characterized by the existence of consumer progression in product usage. We define progression as a specific-case of state-dependence, where consumers need to consume current content or obtain experience with the product today in order to increase future enjoyment with the product. Examples of progression are common and widespread. For example, in role-playing and online multi-player computer, players are required to perform certain tasks that allow them to advance to more challenging (and more rewarding) content. Several popular mobile apps, such as the well-known game “Angry Birds”, have adopted a similar structure to develop consumer interest, where additional “episodes” are launched at a frequent basis to keep interest high. In the TV broadcasting industry, such as in recent serial shows “24”, “Lost”, and “Heroes”, and more generally in soap operas, viewers are “required” to watch episodes in sequence to fully enjoy their experience since each episode adds to the understanding of what will happen in future episodes. This is a similar framework to educational products, where individuals need to learn basic concepts earlier on to be able to comprehend more complex material later on. It is within this broad framework of *progression* and frequent product experiences and content consumption that this paper develops its contribution to the marketing literature.

Modeling the level of consumer involvement with a product involves several research challenges. First, consumers interested in progressing with product usage are by nature forward-looking because any current usage will impact the enjoyment of future content. Hence, consumer choice will depend simultaneously on how much content has been consumed, the current availability of choices, and future product introductions. Second, consumer involvement is itself hard to measure due to most available consumer choice data sets including records of purchases or transactions but not details about product usage. Third, inherent to most experiential products, there are important social interactions among users, either through word-of-mouth, for instance when consumers discuss the events of a recent TV episode, or possibly driven by the need to cooperate to increase content enjoyment, resulting in the creation of online forums or networks of users, as is frequent in the computer games industry.

We propose and estimate a structural model of dynamic product usage that explains how consumers make decisions in environments where user involvement, progression, and product updates are essential components. The model allows for daily consumption decisions, with forward-looking behavior about new content that the firm launches to maintain or increase interest. Our unique data set captures content consumption and the level of involvement of consumers over time. Content consumption is modeled as a function of past participation and experience, availability of current content, and expectations about the introduction and quality of future content. This set up creates path dependency because past decisions of each individual change his or her current choices, leading to the endogenous formation of preferences and a distribution of individual consumer states that influences the evolution of participation decisions over time. Additionally, consumers are modeled as heterogenous in their preferences for content, which is an essential component in explaining decisions regarding experiential products (Holbrook and Hirschman, 1982). Finally, our model also accounts for the effect of social interactions of consumers.

The approach is applied to the online computer games industry, specifically to the popular computer multi-player game World of Warcraft. The game has grown to be one of the best-sellers of all time, with more than 10 million subscribers in 2010, and it provides one of the best examples of recent trends in the online software industry, which include social games made available in platforms like Facebook and new applications for mobile consumption. Additionally, this industry is particularly relevant to measure the relation between innovation and consumer involvement because product updates are frequently introduced to maintain interest and provide incentives for consumer participation.

Our results fall into two areas: explaining consumer participation and providing insights on product innovation. On the demand side, we find significant heterogeneity in consumer preferences, with a segment of about 30% of the market being composed by "hard core" players, while the remaining 70% is divided into two groups of more casual players. The "hard core" early adopter segment consumes content at a much faster rate and their level of participation declines faster than that of the other groups. The casual group of consumers prefers slower and sporadic innovation and can be considered a laggard segment. We also find that the hardcore players have a leading impact on the user community and affect the participation of the other segments. Regarding innovation, we test different scheduling of product updates. We find that keeping a steady launch of product updates is preferable to concentrating innovation at earlier stages of the product lifecycle. The aggregate response rate to innovation is driven by a balance between keeping the multiple consumer segments satisfied and mainly avoiding the exit of early adopters from the market. Finally, we use our model to quantify the value of a product update, measured in terms of consumer participation.

The remainder of the paper is structured as follows. The next section discusses the relevant literature. The description of the model is included in section 3. Section 4 provides details about the data set used in the paper. The estimation algorithm is presented in section 5 and the results are analyzed in section 6. Section 7 describes managerial applications and section 8 concludes.

2 Literature Review

Our work is closely related to three streams of literature. In the modeling side, our paper draws from recent advances in the estimation of dynamic demand models, while in the substantive side, our empirical application is closer to papers on experiential products, the entertainment industry, and the timing of product updates.

Regarding dynamic models, there have been several recent papers that explained demand when forward-looking consumers anticipate a future firm's action. For example, Song and Chintagunta (2003) propose a demand model where consumers are heterogenous in their product and price preferences and forward-looking in terms of price and quality levels. They apply their approach to the camera category and find that an expected future drop in prices leads to the postponement of purchase decisions. Chevalier and Goolsbee (2009) explore the forward-looking behavior of consumers in the textbook category. They find that consumers have correct expectations about the possibility of selling a used book and that they take this fact into account at the moment of purchase, showing that buyers are sufficiently forward-looking so that accelerating the publication of new editions does not increase revenues for publishers. Misra and Nair (2011) study the performance of salespeople, who are forward-looking due to achieving sales objectives of quotas and ratcheting. Salespeople face a future goal, usually at the end of the quarter, and the current level of effort is influenced not only by current compensation, but also expectations about meeting the future objective. In the computer hardware industry, Gordon (2010) proposes and estimates a model of adoption and replacement in the microprocessor industry, where consumers have expectations about future prices and quality of micro-processors. Conlon (2010) uses constraint optimization, similar to our estimation approach, to estimate a dynamic model of demand and supply of LCD TVs. The forward-looking behavior framework is also present in our paper, although motivated in a different way. In our case, players choose to consume content due the possibility of obtaining increased utility from additional product involvement since the firm frequently launches new updates that require some level of past content consumption.

Regarding literature related to experiential products, Holbrook and Hirschman (1982) motivate and describe the main concepts relevant to the study of such products: the definition of task consumption more than product transaction; the inclusion of unobserved heterogeneity to account for user differences in product enjoyment and in the opportunity cost of time; the measurement of the degree of involvement with the product; and learning via a feedback loop of enjoyment, which was also present in early models of buyer behavior (Howard and Sheth, 1969). These concepts were developed in other papers on the entertainment industry, where experiential products abound. For example, past literature has modeled optimal program scheduling (Danaher and Mawhinney 2001, Goettler and Shachar 2001), investigated sources of viewing persistence (Anand and Shachar 2004), and estimated preference interdependence among groups of viewers (Yang, Narayan, and Assael 2006). These three components - scheduling, persistence, and consumer interdependence - are essentials parts of our model.

Finally, there are papers that develop the forward-looking framework in categories involving experiential products. Namely, in the computer software industry, Nair (2007) looks at the impact of forward-looking behavior on firms profits, when consumers are likely to postpone purchase if they expect prices of console video games to go down. The author uses the demand model to compute the sequence of optimal prices that take into account consumer expectations. Ishihara (2011) develops a dynamic demand model that quantifies the impact of the used goods market of video games on prices and profits of firms. Rao (2011) models demand for online purchase and rental of DVDs, where the purchase decision is modeled as a function of individual expectations about the frequency of viewership of each DVD, while Hartmann and Viard (2008) study the effects of future reward on the behavior and loyalty of golf players. These papers focus primarily on the effect of changes in marketing mix variables, such as price and quality, on the purchase behavior of consumers. Similarly, we allow for consumers to have expectations about the timing and quality of future product introductions but our main interest lies on measuring continuous consumer participation and product involvement in the presence of innovation.

3 Industry and Data

Although our approach can be applied to a variety of products, we focus on the online gaming industry as an application of the proposed methodology. In this section, we discuss relevant details about this category. We then describe the data and the product used in our empirical section of the paper.

3.1 Online Games Industry

Individuals are playing video games at a fast growing rate. In the United States, about 40% of adults are now regular players, a number that goes up to 83% when looking at the teenager group. With online connectivity and the ubiquitous presence of the Internet, online or networked games have been growing exponentially, with a recent study showing that about 67% of teenagers regularly play some game online (Rideout, Roberts, and Foehr, 2005), resulting in a worldwide market for online games that was expected to surpass $15 billion in 2010 with additional sales of virtual goods likely to exceed $1 billion (Playlogic Entertainment Inc, 2010). With the development of mobile content, apps, and online gaming console interactions, individuals can now use a number of devices to access both games and social communities, through cell phones, computers and their gaming consoles (Williams, Yee, and Caplan, 2008).

In our application, we use data from the popular online game, World of Warcraft, from the company Blizzard Entertainment, a division of Activision Blizzard. According to the game's website, World of Warcraft is a "Massively Multiplayer Online Role-Playing Game (MMORPG), set in the high-fantasy universe centered around persistent online personae". In other words, the game involves characters that players improve over time and use to explore the environment developed by programmers. The game was originally introduced in 2004, and it was the bestselling PC game of 2005 and 2006 worldwide. Since its introduction, Blizzard has launched three full-fledged expansions and dozens of regular patches that added new content. Each of the three expansions introduced large quantities of new content and in most cases a new chapter in the storyline behind the game, while small product updates or patches either modify previous versions of the game or introduce smaller amounts of content. In 2008, the game had more than 11.5 million subscribers worldwide. The game expansion covered in our data sold more than 4 million copies in the first month alone

(Blizzard Entertainment, 2008).[1]

The game has several components. First, it has a background story that players can progress through by doing tasks in the game (e.g., defeating an enemy). With the completion of tasks, the player's character gains useful experience and statistics that allows it to proceed deeper in the storyline and successfully perform harder tasks. Since the launch of the second expansion, players are able to track their past progress through a measure called achievements. Achievements record the successful first-time completion of a specific (and arguably important) task in the game, as well as the date of completion. These achievements have a social component, since they are public information and players can use them within the player community to show evidence of past performance. Second, the tasks in the game can be divided into individual or group tasks. If a group is required to perform a task, players must currently team up in sets of 5, 10, or 25, to be able to accomplish the task. The difficulty of the task is appropriate for the team size and tasks that require a larger group of individuals tend to require higher expertise from players in terms of game knowledge and ability. Third, although there is content to be enjoyed individually, the game can only be played if the players are logged in into the firm's server. Thus, most actions of the players are visible to others because players interact in the shared environment and are able to influence actions of other individuals.

Blizzard Entertainment offers statistics about the participation and performance of each player in order to provide a more complete experience for users and allowing them to track their progress in the game. Several game-related websites process this information into databases that allow comparisons across players and provide recommendations or advice on how to progress in the game. In this paper, we use a publicly available data set collected from such a site, www.wowhead.com. Additionally, we use information from the official game website about product updates, their content, Blizzard's actions, and announcements. We next provide more details these data.

3.2 Consumer Progression and Product Updates

Our data set includes information about the game from November 13^{th}, 2008, to December 7^{th}, 2010. The time periods included in our data cover the lifecycle of the second expansion of the game in its entirety, from its introduction to the day before the introduction of a new expansion. During

[1] For more information, visit http://us.battle.net/wow/en/.

these time periods, the firm introduced three product updates or patches that added content to the initial expansion launch. With this additional content, more achievements were made available for players, increasing the variety of tasks in the game.

The achievements at the introductory time period and for additional patches are presented in Table 1. Before the introduction of the expansion, players had available content from previous versions of the game, denoted by patch 0 and available for all time periods of our analysis. Patch 1 represents the initial content of the expansion, launched on November 13, 2008, and contained 156 achievements to be performed. After about six months, a new content patch was launched with 143 more tasks. The content of the expansion was completed with two additional patches launched 295 days and 422 days after the introduction, with 23 and 50 tasks respectively.[2]

For our application and to reduce the computational burden, the tasks in each patch are coded into groups depending on the difficulty level and requirements of the tasks. For each patch, we create four groups of tasks, as shown in the fourth column of Table 1. In each patch, the lowest level tasks are individual tasks, easier to perform. Tasks coded as levels two, three, and four, demand more knowledge about the game and the cooperation of multiple individuals to obtain the achievement, respectively 5, 10, or 25 players. Across patches, the level of the tasks continues sequentially, that is, after levels one through four in patch 1, tasks are coded as five through eight in patch 2, and so on. Any deviation from these coding rules for each task is based on game forum discussions and statistics about the difficulty of tasks available in game related websites.[3] With the introduction of each patch, the set of available choice alternatives expands. The resulting choice sets are shown in the last column in Table 1.

3.3 Player Participation and Progression

We have information about the choices of 350 players from one of the game servers. These players were randomly selected from all players who had achieved the required level to play the content made available by the expansion covered in our data by December 2010. We chose not to study

[2]Given the objective of the paper, the data includes achievements related to the game storyline representing progression. There are other achievements related for example to environment exploration that we do not include in our analysis.

[3]We note that it is possible to make these levels a function of task characteristics, so that the increments between tasks is not linear. For example, in the game, the firm gives a level of difficulty/reward to each task. We tested our approach using information specifically related to the level of task assigned by the firm, instead of the coding described in this section, and found no significant differences in our results.

Patch	Release time (in days)	# Tasks in each patch	Task level, coded within patch	Task level, coded overall	Available tasks
0	before $t = 0$	68	1-4	1-4	0, 1-4
1	$t = 1$	156	1-4	5-8	0, 1-8
2	$t = 183$	143	1-4	9-12	0, 1-12
3	$t = 295$	23	1-4	13-16	0, 1-16
4	$t = 422$	50	1-4	17-20	0, 1-20

Table 1: Patches and respective tasks.

the behavior of new players because most of the content described above and included in our data set was dedicated to providing incentives for further participation and progression of experienced players. Hence, the results in our application should be seen as an analysis of the behavior of this segment of consumers and not of the full community of players in World of Warcraft.

For each individual, we observe the date when one of the tasks available in the game is performed. Across individuals, players completed an average of 44 tasks out of a total of 440 tasks in our data, with a standard deviation of 30. On average, 39 tasks are performed daily by our sample of players. The evolution of the daily total number of tasks for the 770 days in our data set and the actions aggregate by product update are shown in Figure 1. Looking at the aggregated pattern displayed in the top left panel, we see the typical shape of a product lifecycle from introduction, growth, maturity, and decline. After the launch of each patch, we observe an increase in participation, caused by the availability of new choices and excitement with new experiences and progress in the storyline. After several weeks, the interest with each patch dies down likely due to satiation and aging of content.

In terms of progression, Figure 2 presents the evolution of player levels, measured as the highest level of any task played until a certain time period. We show the percentage of players at each level for the last week of each patch. For example, the "1" in the horizontal axis represents the last day of patch 1, just before the introduction of patch 2. The figure also displays the volume transition between levels, coded by the respective line color, where a darker (lighter) line represents more (less) players moving from level to level.[4] We observe that there is significant player progression over time with only 8% players remaining at level 6 or below (out of 20 levels) at the end of our data and about 21% succeeding at doing at least one of the highest level tasks.

[4] We note that our data is much more detailed than the patterns presented in the figure, since includes the transition of players at all time periods between each patch introduction at the daily level.

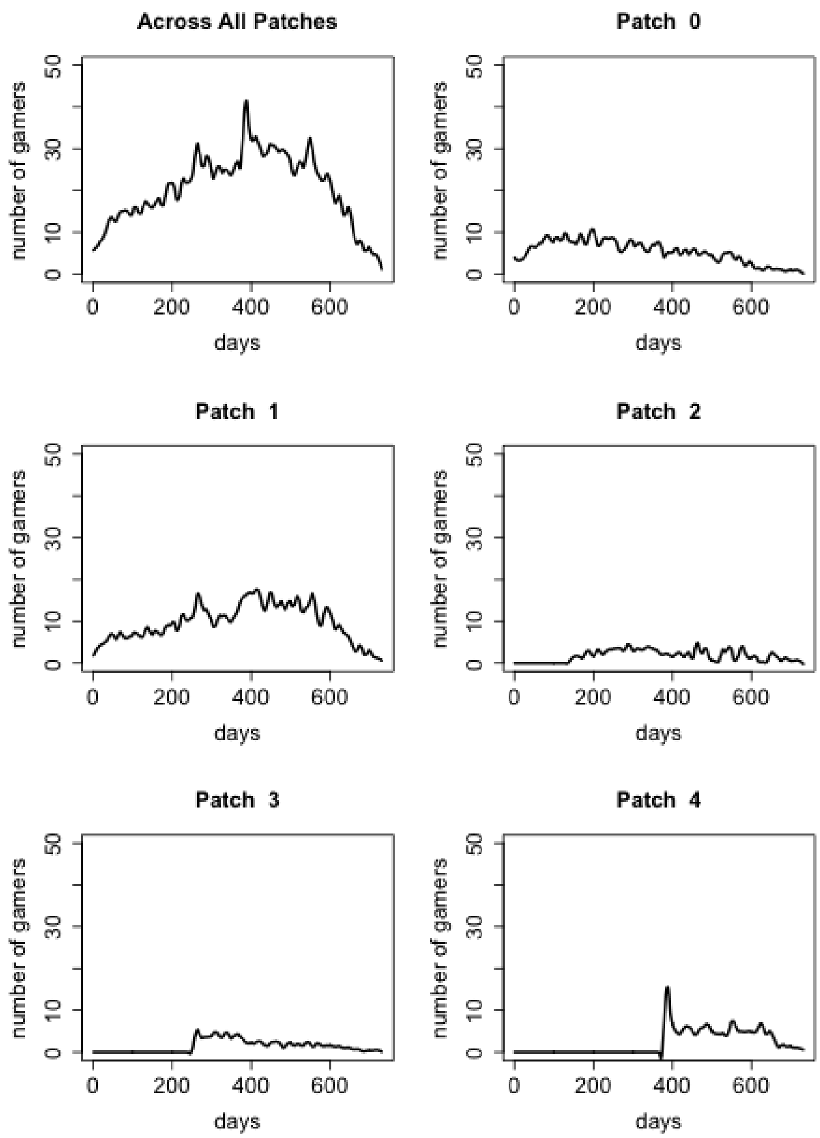

Figure 1: Evolution of the total and by patch number of tasks performed daily by players in the sample.

We also see a significant percentage of players that do not get to successfully complete achievement of the later patches. About 54% of players reach the end of the expansion without being successful at any of the achievement from the last two patches. These could be due to several reasons. First, it is possible that these players have abandon the game during an early stage. Second, these player have a slow pace of progression, and thus are still enjoying previous patch content, and do not require more content to be made available. In our empirical section, we test different scheduling of product updates that influence the participation rate and progression of players.

Moreover, the figure shows that there is significant state-dependence within player progression relative to past achieved levels. For example, we observe that most players at the lowest level at the end of patch one are likely to stay at that level or advance to immediate levels above. On the contrary, players in higher levels tend to advance to the highest level possible in the following product update. This is a result of the design of the game and in general of other products involving progression, where although is is possible to skip content, more enjoyment comes when consumers enjoy content in the order recommended by the the firm. As described in the introduction, this motivation to consume content today to be able to better enjoy additional content in the future is one of the main drivers of forward-looking behavior in our application.

4 Model

In this section, we model consumer choices about content consumption in environments characterized with progression and product updates. We start by outlining our model framework and defining the per-period utility. We then specify the forward-looking behavior of consumers by elaborating on the relevant state variables, their transition, and respective consumer expectations. We then combine the per-period utility with the continuation value of future periods to obtain the decision-making utility that explains consumer decisions. We end this section by describing two additional model details applicable to the case of online games.

4.1 Framework

Consumers draw utility from consuming content offered by a firm. Over time, the firm makes more content available to consumers with the introduction of product updates that increase the choice

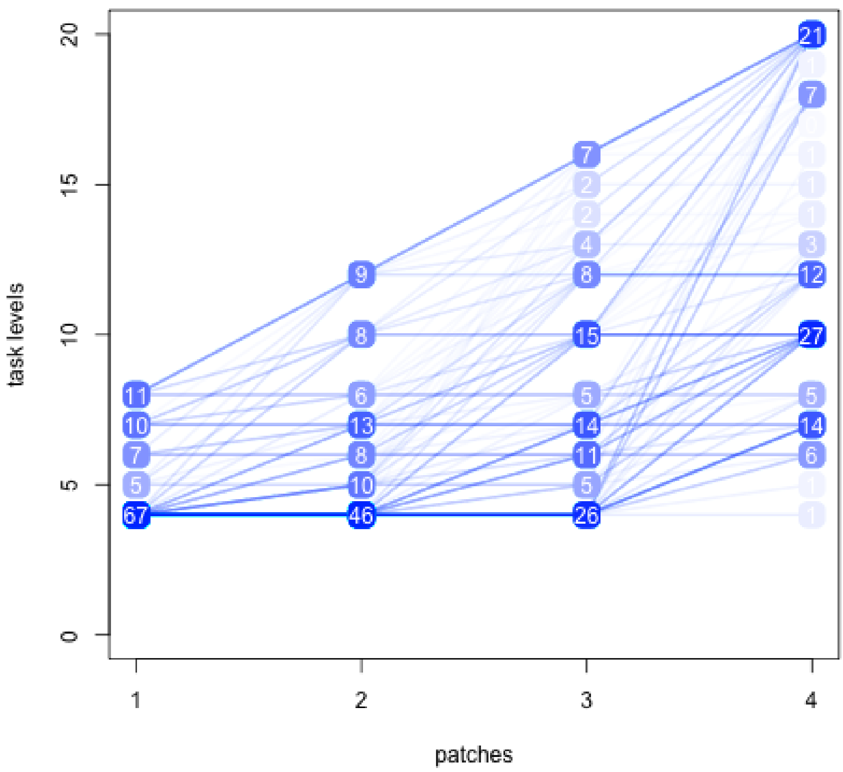

Figure 2: Level of the players and their progression in percentage of the total sample for the last day of each product update.

set. During the product lifecycle, we assume that there are P product updates ($p = 1...P$) with each product update including several alternative content choices.

At each time t, consumers face a set of available content units, $j_{pt} = 1, ..., J_t$, with the subscript p referring to the product update p that introduced choice j_{pt} in the market. Content alternatives are ordered and assigned a level, l_j, with levels varying from 1 to L. For instance, the level l_j can be perceived as the level of difficulty of accomplishing a certain task, as in the case of computer games, or the position in a sequence of episodes or seasons, as in the case of a TV series. Each available content choice can be consumed repeatedly and even though choices are ordered, no particular sequence is enforced to consumers and they can decide to enjoy any of the available units of content or stay out of the market.

The decision to consume content today provides both current and future value. The current utility comes from the per-period enjoyment of using content, while future utility results from the impact of current usage on the utility of future content. Hence, individual decisions are modeled as a sum of two components, the current utility and the discounted flow of future utilities, given a set of state variables.

4.2 Per-period Utility of Content Consumption

Given the set of available alternatives at time t, individual i decides which content unit j_p to consume based on the utility function

$$u_{ij_pt} = \underbrace{\alpha_i + \alpha_j + \alpha_0 X_t}_{\text{Fixed Effects}} + \underbrace{f\left(\beta_{ip}, \tau_{pt}\right)}_{\text{Appeal Time Variation}} + \underbrace{g\left(\gamma_i, l_j, l_{it}\right)}_{\text{Distance to task}} + \underbrace{h\left(\delta_i, l_{it}, \bar{l}_t\right)}_{\text{Community Effect}} + \varepsilon_{ijt}. \quad (1)$$

This utility function captures the main components that influence the consumption decision that we now describe in more detail. In regard to fixed-effects, individual i has some propensity to use content in general and alternative j specifically. These two elements are measured by a heterogenous intercept α_i and a content-specific intercept α_j. Additionally, X_t is a matrix of dummy variables that capture time-specific characteristics that influence patterns of content consumption, such as holidays or weekends.

Over time, content may gain or lose appeal due to consumer excitement build-up or satiation with the product. We allow for the appeal of choice j_p to vary as a function of time through the term $f(\tau_{pt})$, where τ_{pt} is the number of days from the introductory period of update p until period t. The function $f(.)$ models the effect of the time on the consumption utility and after testing a number of alternatives we chose the quadratic function to capture possible non-linearities

$$f(\beta_{ip}, \tau_{pt}) = \beta_{1ip}\tau_{pt} + \beta_{2ip}\tau_{pt}^2. \tag{2}$$

As described previously, progression plays an important role in this type of products. Consumers are likely to be attracted to content that suits their product involvement resulting from past content consumption. The model captures this through the match between the state of the player in terms of content involvement at time t and the content level l_j. To do so, we define l_{it} as the value of the highest level of any content enjoyed before time t by individual i. The utility of each task j is influenced by the distance between its level and that of the consumer and we allow for coefficients to vary if individual i is above or below the task level,

$$g(\gamma_i, l_j, l_{it}) = \begin{cases} \gamma_{1i}(l_{it} - l_j) & if \quad l_{it} > l_j \\ \gamma_{2i}(l_{it} - l_j) & if \quad l_{it} \leq l_j \end{cases}. \tag{3}$$

We expect that alternatives of levels neighboring the level of the user to be the most enjoyable with a good match between player and content levels.

A third component of the model deals with the social aspect of content.[5] We model this effect by comparing individual i's level l_{it} and the mean level of other consumers at time t, $\bar{l}_t$. If there is a social prestige effect then we expect that individuals prefer levels higher than the community level. We define D_{it} as a indicator variable that takes the value of 1 if $l_{it} > \bar{l}_t$ and zero otherwise and set the community effect as a change in the utility of being above the community mean, $h\left(\delta_i, l_{it}, \bar{l}_t\right) = \lambda_i D_{it}$.[6]

We define consumer heterogeneity as a finite mixture of segments. With this approach, the

[5] In most cases, online games have a social component where players need to interact with each other to progress, for example, to do tasks together. The social component is also present in the case of TV shows, since viewers tend to discuss recent developments in the storyline of TV shows with friends.

[6] Although other functions are possible, such as using the distance between the level of the player and the level of the community, our choice of function $h(.)$ is driven mostly by the computational requirements of the dynamic model and the number of states required to model more complex social effects.

unconditional probability of individual i choosing j will be a result of a product of two components. First, based on the history of choices, each consumer i has a probability to be assigned to segment k. Then, conditional on belonging to segment h, consumer i will have a probability of consuming content j. After testing different number of segments, we chose to estimate our model with $K = 3$ segments, based on fit, number of parameters, and computational demands of the model. Instead of having individual parameters, we will obtain segment-specific parameters. We note that there is additional heterogeneity in our model because consumer level l_i is individual specific.

Finally, we assume that ε_{ijt} are shocks that are independently and identically distributed across individuals, time, and content, with a Type I extreme value distribution. The non-random per-period utility of choosing the outside alternative is fixed at zero for identification purposes. The parameter set to be estimated is denoted by $\Theta = \{\alpha, \beta, \gamma, \delta\}$.

4.3 State Transitions and Consumer Expectations

At any period t, the variables that influence the state of consumer i include: 1) individual level, l_i; 2) the relative position of l_i compared to community level, $\bar{l}_t$, captured by D_{it}; 3) the index of the most recent content batch, denoted by $\tilde{p}_t$, which defines the choice set available to consumers; 4) the number of time periods that have passed since the introduction of a content batch p, τ_{pt}; and 5) time-specific effects, X_t. We collect these state variables into the set $S_{it} = \{l_{it}, \tau_{pt}, X_t, D_{it}, \tilde{p}_t\}$. We start by discussing the state variables that have deterministic transitions, l_{it}, τ_{pt}, and X_t, and continue by describing the expectations about the transition of state variables where there is some uncertainty, $\tilde{p}_t$ and $\bar{l}_t$. We conclude this section by discussing expectations about the quality of future content captured by α_j.

The first state variable, which measures the level of the consumer l_{it}, takes the value of the highest level of any content enjoyed in the past by individual i. Individual i moves to a new level that is equal to the chosen content level if this content's level is higher than consumer i's current level l_{it}. Formally, the state transition of this variable is deterministic from the individual's point of view because the choice of content is a control variable[7] and its transition rule is given by

$$l_{it+1}\left|l_{it}\right. = max(l_{it}, l_{ait}), \tag{4}$$

where l_{ait} is the level of the chosen content a.

The state variable τ_{pt} indicates the number of time periods that have passed since the introduction of the last product update. This variable captures the age of each available unit of content, in days, since its introduction and it evolves as $\tau_{pt+1}\left|\tau_{pt}\right. = \tau_{pt} + 1$, with $\tau_{pt} = 0$ for any t before the introduction of content patch p. Finally, the time effects X_{jt}, which include several dummy variables for weekend, christmas, thanksgiving season, etc., also evolve deterministically. In our application, most of these time effects are changes in content that match the respective season and consumers know with certainty when these time effects occur because they appear with the same frequency and on the same days of each year.

With respect to the variables with uncertainty and concerning product updates, the index of the most recent content update $\tilde{p}_t$ at time t defines the set of tasks available to consumer J_t. The introduction of a new product update adds several alternatives to the existing choice set but consumers are uncertain about the timing of the new content release. The probability of new content release grows with the time since the last introduction and we assume that consumers start anticipating the release of new content after m days from the most recent release. The probability of seeing new content release is defined as

$$Pr\left(\tilde{p}_{t+1} = \tilde{p}_t + 1\left|\tilde{p}, \tau_{pt}\right.\right) = \begin{cases} \phi\left(\tau_{pt}\right) & if \quad \tau_{pt} > m \\ 0 & otherwise \end{cases}, \tag{5}$$

while the probability of not observing a product update is given by $1 - Pr\left(\tilde{p}_{t+1} = \tilde{p}_t + 1\left|\tilde{p}, \tau_{pt}\right.\right)$. The probability $\phi\left(\tau_{pt}\right)$ is obtained based on the empirical distribution of previous launches of product updates observed in the data.[8]

Furthermore, at each time period t, we assume that consumers have true expectations about

[7] Later in this section, we discuss the case where there may be some uncertainty about the outcome of an action, which is possible in online games when a gamer attempts an activity but is not successful.

[8] In the case of our empirical application, the probability of seeing new content release has the additional condition that it is only positive on Tuesdays, when server maintenance is done and updates are made available to the public. The choice of m is based on two consumer information sources: discussion forums about the game and the schedule of past announcements by the firm about future product updates.

D_{it+1}, their relative position when compared to the community mean level in the next time period. Consumers develop such expectations by observing their current relative position, their choice of action, and other variables that might influence the evolution of the levels, such as a recent introduction of a new product update. We estimate the probability of consumes expecting to be above or below the community level with a logit model where the dependent variable is an indicator of consumer's position against the mean community level at $t+1$, and the explanatory variables are the above mentioned variables at t. The logit model is estimated with observed data before the estimation of our model and based on the obtained coefficients, we use the value of the explanatory variables to obtain probabilities of seeing a 1 or 0 for D_{it+1} during the estimation. The transition rules for all variables is denoted by $\pi(S_{it+1}|S_{it})$ in the remainder of the paper.

4.4 Consumer Choice

At the initial period consumers observe the launch of content and decide whether to do any of the available actions j during the following T periods. Consumers are assumed to know T, that is, how long content will be available based on other similar products or on previous versions of the content. Hence, we state the consumer dynamic problem as a problem with the finite horizon.[9]

Based on the previous sections, we develop the choice probabilities of consumer i. At time t, each consumer maximizes the expected present value of content utility by choosing the optimal sequence of actions a such that

$$V_t\left(S_{it}, e_{it}\right) \equiv \max_{\mathbf{a}} E_{\mathbf{a}} \left(\sum_{t'=t}^{T} \delta^{t'-t} \left[u\left(S_{it'}, a_{t'}\right) + e_{ia_{t'}} | S_{it}, e_{it} \right] \right). \tag{6}$$

Using the Bellman principle of optimality and the assumption about the Type I extreme value distribution of the shocks, consumer i's value function conditional on choosing action a at time t can be written as (net of e_{iat}):

$$v(a_{it}, S_{it}; \Theta_i) = u\left(a_{it}, S_{it}; \Theta_i\right) + \delta \int_S \bar{V}_{t+1}(S_{it+1}; \Theta_i) \pi(S_{it+1}|S_{it}, a_{it}) dS_{it+1}, \tag{7}$$

In equation 7, a_{it} is the agent's content choice, $\pi\left(.\right)$ denotes the transition probabilities between

[9]It is possible to make T a function of the time since the last product update. The problem is still a finite-horizon one but with a flexible terminal time period.

states from time t and $t+1$, and $\bar{V}_{t+1}(S_{it+1};\Theta_i) = log\left(\sum_{a'\in J_{t+1}} exp\left[v(a'_{it+1}, S_{it+1};\Theta_i)\right]\right)$ is the value function integrated over the extreme value distributed shock to consumer utility.

In the last period T, consumers get a final choice utility u_{ijT} and a termination value C_{iT} , which we make a function of the consumer's level l_{iT}:

$$v(a_{it}, S_{iT};\Theta_i) = u\left(a_{iT}, S_{iT};\Theta_i\right) + C_{iT}\left(l_{iT}\right). \tag{8}$$

Given the assumption about the extreme value distribution of the random shock to utility, the probability of choosing action j at time t by individual i becomes

$$Pr(a_{it}, S_{it};\Theta_i) = \frac{exp\left(v(a_{it}, S_{it};\Theta_i)\right)}{\sum_{a'\epsilon A_t} exp\left(v(a'_{it}, S_{it};\Theta_i)\right)}. \tag{9}$$

These individual probabilities will be used directly in the estimation routine to obtain parameters in Θ_i, since we observe the consumer action a_{it} at each time period t and states S_{it}.

4.5 Modeling Details for the Online Games Environment

There are two additional complexities that arise in the particular case of online games: (1) the possibility of failure at a chosen task and (2) the uncertainty about the quality of future choices α_{ij}. We now describe in more detail how we address these issues.

First, in most computer games, there is a chance of failure when a player attempts a task. The uncertainty of success is a challenging component of games and part of the incentive for players in competitive environments. In the case of our application, we assume that an individual is aware that if a task is chosen, especially a harder task, it is possible that he will not be successful and hence he will not move up to the intended task level. To account for this possibility, we use data on success rates for each task and player level[10] to build a matrix W, of dimensions $[L \times L]$, where $w_{l_i l_j}$ is the element of matrix W that measures the rate of a player of level l_i succeeding at performing a task of level l_j. The main difference between a failure and a success of attempting content j comes from the fact that, although the players still obtain a per-period utility of an attempt, succeeding at the action of a higher level leads to progression, while failing at accomplishing a task will make

[10]The data was obtained at worldoflogs.com and success rates were averaged across achievements within a level.

a consumer remain at its current level. Hence, the continuation valuation will be different between the two outcomes and it will influence individual choices.

In addition, we also need to correct the choice probabilities to take into account the success rates. The probability of seeing a positive outcome of doing action j is a result of the product of two probabilities: the probability of consumer i choosing task j and the probability of the consumer succeeding in that attempt. Thus, using equation 9, the choice probabilities $\hat{Pr}_s(a_{it}, S_{it}; \Theta_i)$ accounting for the possibility of failure that enter the likelihood function take the following form:

$$\hat{Pr}_s(a_{it}, S_{it}; \Theta_i) = Pr(a_{it}, S_{it}; \Theta_i) \times w_{l_i l_j}. \tag{10}$$

Second, we discuss the quality of tasks to be introduced with future product updates. With the future introduction of new product updates, more tasks will be offered by the firm and the intrinsic utility of each of these future tasks is captured by the parameter α_j. Before the availability of these choice alternatives, it is possible that players are unsure about the values of α_j and they form expectations, which we denote by $E(\alpha_j)$. We assume that consumers form these expectations based on observed past launched content. Thus, before the launch, $E\left(\alpha_{k_p}\right) = \alpha_{k'_{p-1}}$, and become a realized $\alpha_{k_{p+1}}$ after the launch of patch $P+1$.[11]

5 Estimation

In this section, we discuss some limitations of the data and the respective assumptions, identification of parameters, and the estimation algorithm.

5.1 Data Limitations and Assumptions

Our data has several limitations that we want to point out. First, the main limitation of the data regards the fact that we only observe the first time that consumers perform an action. Thus it is possible that consumers participate and repeat the same action, instead of doing an alternative task for the first time. This limitation is less severe in our modeling approach due to the fact that we group a large number of tasks into a smaller set of alternatives that consumers can choose to

[11] We also tested using observed characteristics of past content alternatives to form expectations about the quality of future content with insignificantly different results.

perform. Hence, even though consumers may repeat previously obtained achievements, we will see them participate if they do at least one task in the same day that has not been done before from the same group of tasks. Additionally, in our case, we are interested in modeling progression and new content adoption and so the decision of consumers in our application can be viewed as the decision to consume new content and the outside good will capture repeat content consumption.[12]

Additionally, we do not observe consumer payment of subscription fees. Hence, our analysis and counterfactuals presented in the results section will not be able to provide exact predictions of revenues. We can however obtain estimates of average participation rates for individuals and with the assumption that a monthly subscription of $15 is equivalent to a daily contribution of $.50 to the firm's revenue, we are able to provide estimates of the value of a product update and compare revenues of two alternative scenarios.

In terms of assumptions, we want to briefly discuss our choice of a finite horizon for our application. First, the firm previously launched other products with product life cycles very similar to the product studied in our application. Second, the firm announced the date of a posterior very large product update (or expansion) for the product in an early stage of the periods in our analysis. This expansion, although set in the same online environment, did not add to the storyline of the product nor achievements in our application. We use the day before the date of launch of the additional expansion as a terminal time period for the analysis. Our model can be modified to allow for the final date to be flexible, with additional consumer expectations about this terminal timing.

Finally, in terms of discount rate, we present results with $\delta = 0.99$ per day, which corresponds to 0.932 per week and inline with other discount rates used in entertainment or experiential products. For example, Hartmann and Viard (2008) uses a similar value, while Ishihara (2011) uses 0.885 per week for video games. See Yao et al. (2011) for a list of other discount rates choices. We tested other discount rates with very similar substantive results.

[12] In practice, we can implement the following change in the probabilities to account for this limitation: every time a consumer does one achievement included in alternative j, we can move $\frac{1}{C_j} Pr(j, S_{it}; \Theta_i)$ from the probability of chosing action j and add it to the outside good probability, where C_j is the number of achievements included in choice j.

5.2 Identification

We briefly discuss the identification of the different parameters of our model. The content intercepts are identified by the average observed rates of participation for each content alternative. We opt for a segment-specific base intercept to capture the propensity to play the game and task-specific intercepts for each available action in the game common to all individuals, with one task intercept set to zero for identification. The satiation coefficients are identified by the decline in choices to do achievements in each product update over time.

In terms of the parameters related to the distance between individual and task levels, these are identified from the progression rate of individuals. If for example consumers tend to do most tasks closer to their level, then the match between task and individual levels is important to their enjoyment of content and the value of the parameters will reflect this behavior. Additionally, the community effect is identified by the variation in actions of consumers below or above the mean level of the community. Moreover, both individual and community levels are moving over time, which implies that the identity of consumers above and below the mean will also change over time. Hence, variation across individuals and across time identifies the community effect.

Finally, we discuss the nature of the forward-looking behavior captured by the continuation value in the utility function. This effect originates primarily from the variation in choice sets with the introduction of new product updates. With the expectation that more tasks will be available in the future with the introduction of new product updates and that those tasks are of higher levels than current tasks, consumers have additional incentives to continuing progressing in the game. In other words, the general pattern of progression, and more specifically, the patterns of doing higher level tasks, even in the presence of satiation and especially when a new product update is about to be released, reveals the importance of the continuation value on the utility function.

5.3 Estimation Algorithm

We estimate our model using the Expectation-Maximization (EM) algorithm (Arcidiacono and Jones, 2003) and combine this algorithm with the use of constraint optimization. The mathematical program with equilibrium constrains (MPEC) approach (Su and Judd, 2011) helps estimation in the finite horizon case, allowing us not to use backward recursion to obtain the consumer value functions from time T to the initial time period. The use of the MPEC approach also reduces the computational burden of obtaining conditional expectations of the value function for each possible combination of the state S, action a, and time period t.

Given observed choices made by each of N gamers, the log-likelihood of observing the data Y is:

$$LL(Y|\Theta,\Lambda,V)=\sum_{i=1}^{N}\left(\sum_{k=1}^{K}\left[Pr(i\in k)\sum_{t=1}^{T}log\left(Pr(a'_{it},S_{it};\Theta_i\,|i\in k\,)\right)\right]\right), \tag{11}$$

where $Pr(i\in k)$ is probability that individual i belongs to segment k and $Pr(a_{it},S_{it};\Theta_i\,|i\in k\,)$ is the probability of individual i choosing action a'_{it} which was observed to be chosen, conditional on belonging to segment k. The vector Λ collects the discrete segment sizes, $\Lambda=\{\lambda_1,...,\lambda_K\}$, and a_{it} stands for the action chosen in period t.

Using this likelihood function, we define the optimization problem for our model as follows:

$$\max LL\left(Y|\Theta,\Lambda,V\right)\quad subject\,to \tag{12}$$

$$\bar{V}_t(S_{it};\Theta_i)=log\left(\sum_{a'\in A_{t+1}}exp\left[v(a'_{it},S_{it};\Theta_i)\right]\right)\qquad for\,t=1,...,T-1,\;and \tag{13}$$

$$\bar{V}_T(S_{iT};\Theta_i)=C_{iT}\left(l_{iT}\right). \tag{14}$$

The estimation algorithm proceeds as follows:

Step 1. In the first step, we make an initial guess for the prior probability that each individual i belongs to segment k, $Pr^0(i\in k)$.

Step 2. For each iteration m, we compute the segment sizes using:

$$\lambda_k^m=\frac{\sum_{i=1}^{N}Pr^m(i\in k)}{\sum_{i=1}^{N}\sum_{h=1}^{H}Pr^m(i\in k)}. \tag{15}$$

Step 3. With $Pr^m(i\in k)$ fixed, we optimize the log-likelihood function in equation 12, subject to constraints in equations 13 and 14, for parameter set Θ^m.

Step 4. Given the obtained parameters Θ^m, we compute the posterior probability that individual i belongs to segment k using the Bayesian updating formula,

$$Pr^{m+1}(i \in k) = \frac{\lambda_k exp\left(\sum_{t=1}^{T} a_{it} log\left(Pr(a_{it}, S_{it}, \Theta \left| i \in k\right.)\right)\right)}{\sum_{k=1}^{H} \lambda_k exp\left(\sum_{t=1}^{T} a_{it} log\left(Pr(a_{it}, S_{it}, \Theta \left| i \in k\right.)\right)\right)}. \tag{16}$$

Step 5: Repeat Steps 2 to 4, until $\left|LL\left(Y|\Theta^m, \Lambda^m, V^m\right) - LL(Y|\Theta^{m-1}, \Lambda^{m-1}, V^m)\right| < \kappa$, where κ is chosen to be a small constant (in our case, $\kappa = 0.0001$).

6 Results

In this section, we discuss the empirical results. We divide this section into three parts. First, we use some fit statistics to provide evidence that the model is able to explain the data well. Second, we analyze the parameter estimates and discuss respective implications. Finally, we present several counterfactual analyses intended to measure the change in user participation with different scheduling of product updates and when one of the product updates is not introduced.

6.1 Fit Statistics

After obtaining the estimates of the parameters, we predict user actions and compare them to actual choices. Figure 3 shows aggregate statistics over each product update. To obtain the fitted plots, we collect the individual probabilities of the different actions given the model parameters and data and aggregate them over individuals and over tasks of the same patch to obtain participation for each time period and for each patch. The lines are smoothed for clarity. The first panel shows actual and estimated participation over all tasks. The remaining panels show participation for each of the product updates. We see that our model explains well the overall pattern of player participation over time and for each of the product updates in our data.

Additionally, we also compute other fit statistics of our model. In terms of total participation, our hit-rate is 92%, that is, we are able to predict very well the timing when individuals do any available task. The likelihood of our model is significantly better compared to an alternative specification where consumers are myopic and to alternative specifications with just one or two consumer segments

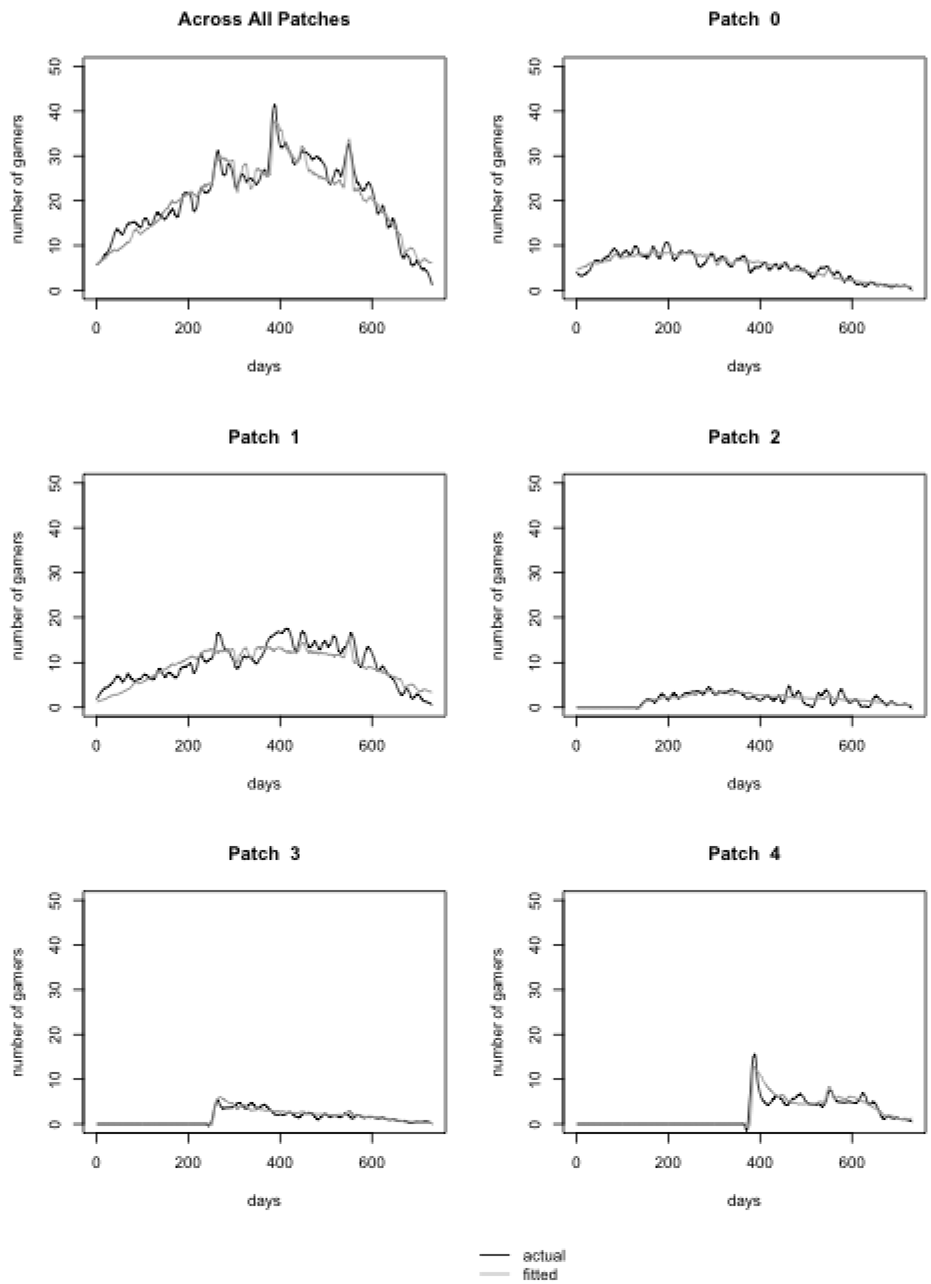

Figure 3: Evolution of the total number of tasks performed per day, over all content and for each product update.

Parameter			Segment 1	Segment 2	Segment 3
Base Intercepts		α	-13.297 (0.375)	-6.835 (0.081)	-8.111 (0.118)
Appeal Variation	Previous content	β_{10}	2.665 (0.111)	1.325 (0.041)	1.586 (0.019)
		β_{20}	-0.190 (0.010)	-0.175 (0.006)	-0.161 (0.003)
	Initial Content	β_{10}	3.010 (0.122)	1.065 (0.601)	1.409 (0.019)
		β_{20}	-0.264 (0.011)	-0.203 (0.011)	-0.174 (0.004)
	Product update 1	β_{11}	3.258 (0.108)	0.111 (0.132)	0.507 (0.071)
		β_{21}	-0.423 (0.019)	-0.192 (0.032)	-0.139 (0.018)
	Product update 2	β_{12}	4.158 (0.460)	-1.851 (0.148)	0.255 (0.081)
		β_{22}	-0.769 (0.020)	0.090 (0.039)	-0.253 (0.023)
	Product update 3	β_{13}	6.739 (0.332)	-2.911 (0.060)	-0.046 (0.128)
		β_{23}	-1.788 (0.804)	0.569 (0.022)	-0.204 (0.729)
Distance to Task	lower levels	γ_1	-0.005 (0.019)	-0.064 (0.004)	-0.003 (0.005)
	higher levels	γ_2	-0.674 (0.013)	-1.131 (0.035)	-0.645 (-0.020)
Community Effect		δ	-0.118 (0.241)	0.580 (0.023)	0.235 (0.058)
Segment size		λ	0.258 (0.010)	0.338 (0.010)	0.404 (0.010)

Table 2: Parameter estimates with standard errors in parenthesis

instead of three segments.

6.2 Model Estimates

Table 2 presents the parameter estimates for the model with three consumer segments, with the standard errors in parenthesis. Looking at the base intercept α, we find significant heterogeneity across segments in the propensity to consume content. Segment 1, with an estimated size of 26%, is composed of players that get the lowest satisfaction from content with a very negative intercept (-13.3), while segment 2, with 33% of players, shows the opposite result, with the least negative intercept (-6.8). Segment 3, with a size of 40% of the market, shows a propensity to play that lies in between the two other segments (-8.11).

Regarding the temporal variation of the content appeal, we evaluate the linear and quadratic terms in the function $f(.)$ for each patch. In order to visualize the temporal variation, figure 4 shows the shape implied by the estimates. We find that Segments 1 and 3, which have a lower propensity for playing, need time to get excited about the product, with a lower participation in earlier stages. On the contrary, consumers is segment 2, who have the highest propensity to play, become satiated faster and participates in earlier stages after the introduction of a product update.

Looking at the coefficients of distance of consumer levels to task level, we first note that the coefficients for tasks of higher levels than the level of a player are significantly more negative than

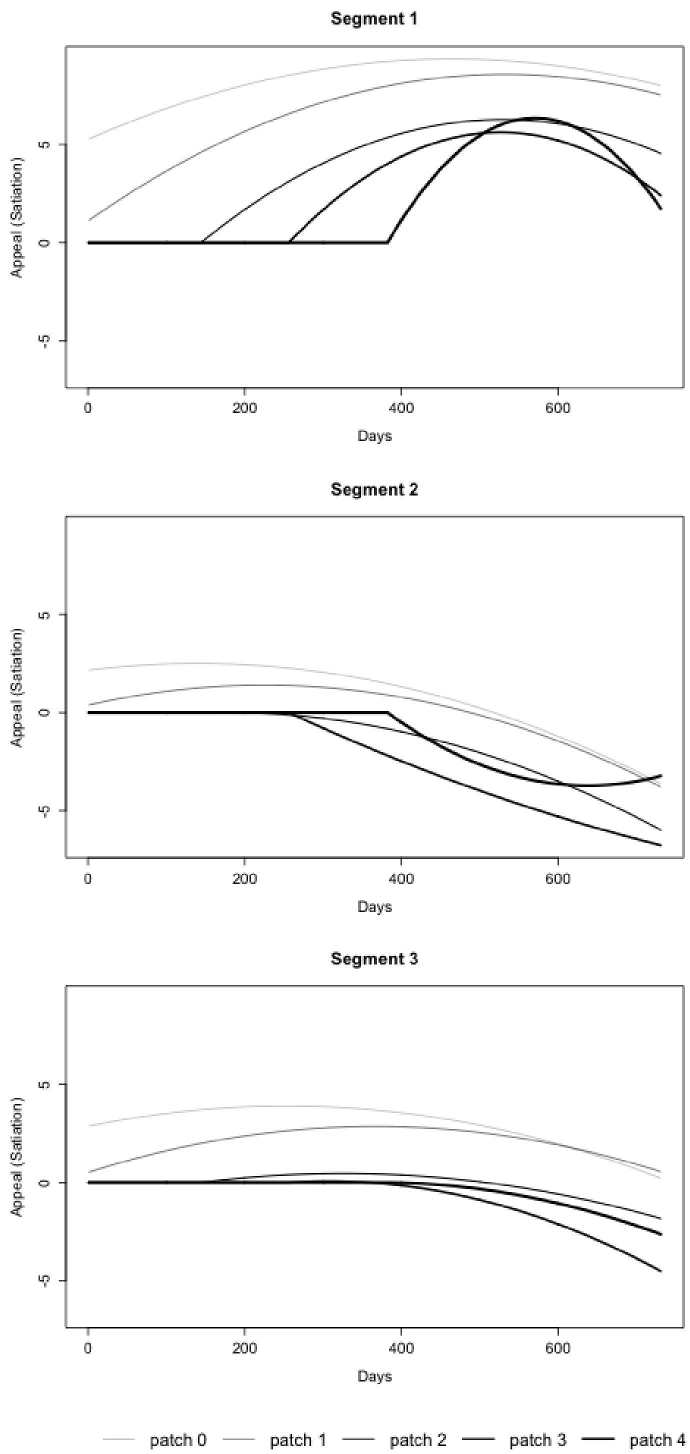

Figure 4: Temporal variation in game appeal, for each patch and segment.

the coefficients for tasks of lower levels. This is a reasonable result because content becomes more challenging with the level of the tasks, either by increased difficulty or by additional required effort in terms of time invested in the game. We also find that individuals in segment 2 prefer tasks that are closer to their level and have higher disutility for more distant tasks. This results implies that these consumers prefer to follow the full progression offered by the game moving upwards one level at the time without skipping levels of content. Moreover and looking at the community effect, segment 2 seems to view the game as a competitive environment and prefers to be ahead of the player community level the most, with a positive coefficient of 0.58. Combined with the higher propensity to play discussed previously, we conclude that this segment includes more serious (hard-core) players, who enjoy progression and want to experience all aspects of the game.

The other two segments show a more casual attitude towards the game in a number of dimensions. The effect of distance between task and individual levels is less impactful in segments 1 and 3, with coefficients closer to zero, leading them to be more likely to skip some content. Additionally, these consumers value less being ahead of the community and segment 1 even prefers being below the average community level. In this game, this can be a beneficial situation because the lower level players can free-ride with more advanced players to complete harder group achievements in the game.

The outcome of these estimates is a clear definition of three distinct segments and we further analyze their behavior using the individual probabilities of belonging to each segment and the observed levels of tasks performed. In the three panels in Figure 5, we show the temporal evolution of the levels of each segment measured as the average level achieved by players in that segment, the observed average mean level of the entire population, and the average levels of tasks done at each period. The players are assigned to the segment with the highest probability $Pr(i \in h)$ given by the model.[13]

We observe that the progression paths are significantly different across segments. Players in segment 1 take almost 400 days to start doing tasks that allow progression and can be considered laggards in terms of content consumption. In fact, for the first 200 periods, these players have virtually no achievements. On the other hand and as previously discussed, segment 2 is composed of more serious players that want to progress in the game. Moreover, although this segment reaches

[13]The probability of belonging to one of the three segments is above 80% for most individuals.

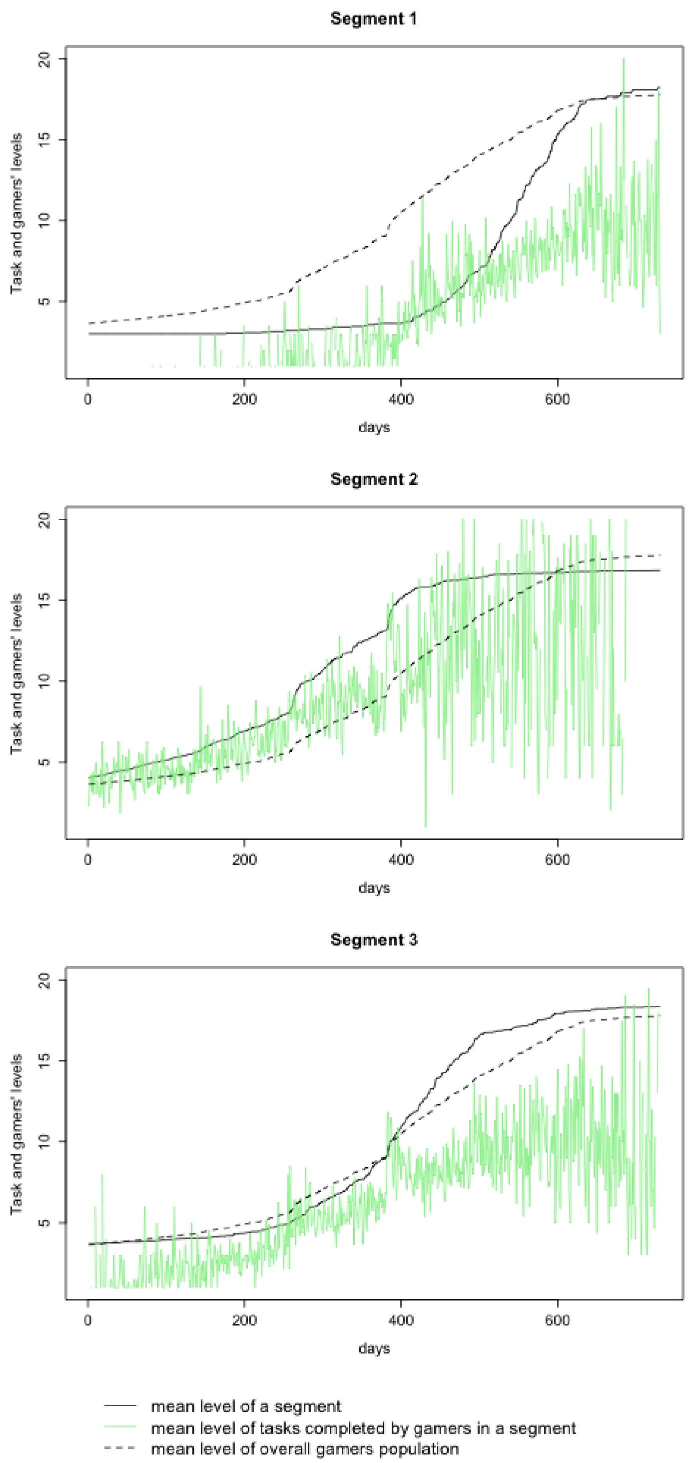

Figure 5: Evolution of the players' level, the mean level of the player community, and the average level of tasks performed in each day.

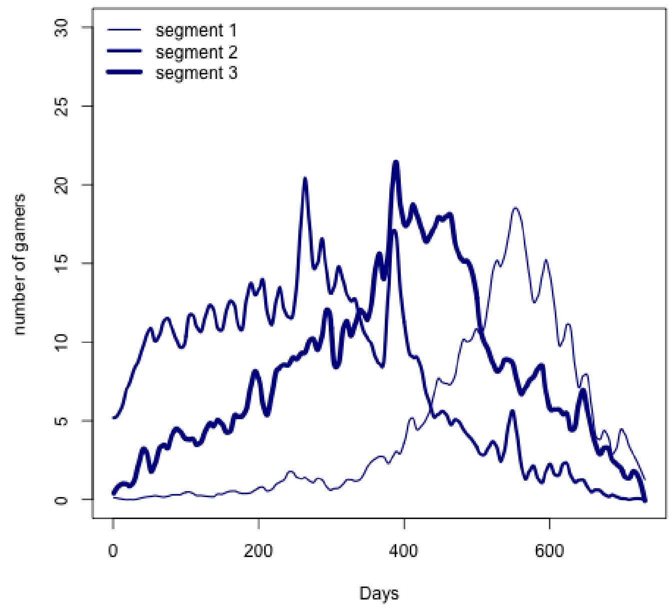

Figure 6: Participation by consumer segment.

a very high level earlier on, their participation does not end. Instead, they continue to consume content of a wider variety of levels, as shown by the average level of tasks performed in each day. Finally, segment 3 is composed of the average player that closely follows the average level of the player community.

We note that there is a large difference between the average level of the segments, from time periods 200 to 500. This is caused by the difference in preferences for the timing of content consumption that emerges from three components of our model: temporal variation of game appeal, distance between task and player levels, and the community effect. The gap between segments shrink at the end of the product lifecycle when the slowest players catch up and the heterogeneity in consumer levels is considerably smaller. Figure 6, where we aggregate content consumption over players by segment, shows the temporal differences more clearly, with segment 2 proving to be the early adopter while segment is the laggard in the market.

6.3 Managerial Implications

We describe two managerial applications of our model. First, we compute the value of one of the product updates. Second, we test a different schedule of innovation by postponing the introduction of one product update. The results provide insights to managerial decisions such as how many product updates should a firm launch over the lifetime of a product and what is the impact of innovation timing on consumer participation and involvement.

6.3.1 Measuring the Value of a Product Update

To measure the value of a product update in our application, we run a counterfactual scenario where we simulate the case where the firm does not launch one of the content patches and compare consumer activity in this scenario with the actual participation. To obtain the simulated participation, we use the estimates of the model and assign a total of 350 individuals to segments in proportion to the estimated segment sizes. We then forward simulate their actions and states based on our proposed model with the difference between the two scenarios being the launch of the third product update. We repeat this exercise 20 times for each of the two scenarios and average content consumption to minimize simulation error resulting from draws of actions from the choice probabilities. We then aggregate participation over individuals to measure content consumption for the player population.

Figure 7 displays the evolution of the total participation over time for the actual and counterfactual scenarios measured as the number of players that do any action j for time period t. We observe that without the third product update, which was launched at period 295 in the actual scenario, participation decreases significantly. This is because consumers do not have an increased choice set from the new content and also due to the effect of satiation with previous patches. This drop has a persistent effect with additional content launched later not making up for the initial loss. This is evident by the reduction in the impact of the last product update (launched at period 422 in the actual scenario) in the counterfactual situation, which is much smaller than in of actual scenario.

To better quantify the impact of this scenario in consumer participation, we computed the difference in the number of actions performed over all time periods and consumers. The total number of tasks observed in the actual scenario was 15,075, while in the counterfactual scenario this number dropped to 13,913, a decline of around 7.5%. The three segments responded differently to the removal of the third update. Segment 2, composed of the more serious players, had a decline of about 10% in their participation rate, while segment 1, composed of the more casual players, only saw a drop of about 4%. Hence, the existence of more content and updates has not only a significant impact on total participation, but also on the composition of consumer types in the game

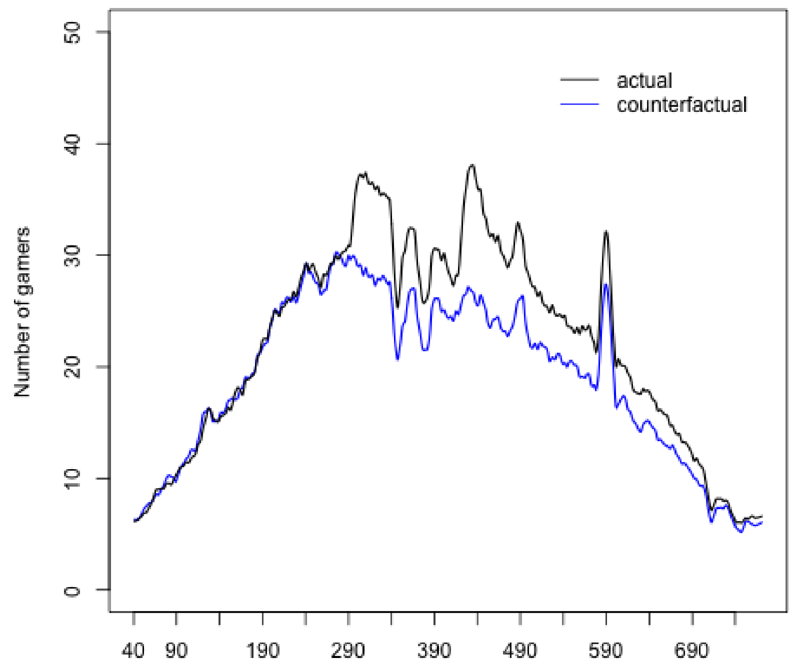

Figure 7: Evolution of player participation from the actual scenario and a scenario where product update (3) is not introduced.

involvement.

Assuming that each player is paying $15 in subscription fees, which is the most common form of payment to participate in the game, the 7.5% decline in participation is equivalent to a the monthly value in revenues of $1.16 per player. There were 18,700 players in the server from which the sample of players was obtained. This implies that for this server alone and at a monthly level, the product update generated about $21,620 in revenues.[14]

6.3.2 Testing Different Schedules of Product Innovation

Our approach can also be used to provide insights about alternative scheduling of product updates. In our data, we observe most updates concentrated in the first half of the product lifecycle, leaving almost 300 days until the end of the product lifecycle without a major product update.[15] In this

[14]It is possible to model the decision to subscribing to the content as a separate decision from consuming content. The purchase decision would take into consideration the expectations of consuming content in future periods and consumers would have an additional level of forward-looking behavior. We do not model this decision because we do not have information about subscription decisions.

[15]The firm did launch a small update with content and tasks at time 590, but it was not related to previous content, had a very small number of tasks, and it was perceived by players as a short-term event. This is reflected in the time series by the spike at time periods 590 to 592 in consumer participation during that event. This is controled for a

counterfactual analysis, we postpone the launch of the fourth product update by two months, from day 421 to day 477. We then use the same approach as in the previous section to obtain estimates of total and segment participation.

Figure 8 shows the comparison between the actual and counterfactual total participation, measured by the number of tasks done by all players at each time t. We see that postponing the last update leads an initial drop due to lack of innovative content. However, this drop is compensated by the increase in participation at the later time period and a persistent positive difference over the last two hundred periods of the lifecycle. Aggregating over time, participation increases by 4% with the postponement of the update. This is justified by two factors. First, consumers were still consuming previous content when the update was introduced in the actual scenario. Thus, this earlier introduction led to increased substitution between tasks and did not attract consumers that would have stayed out of the game. Second, by allowing more time between updates, the overall consumers are better "prepared" to face the content of the last update, leading to a positive and long-term effect on utility. By segment, we see that segment 2 increases participation the most with this new schedule, by 8%, while the casual player segments 1 and 3 increase consumption by 3%.

7 Conclusion and Future Research

In this paper, we propose a model of content consumption by forward-looking consumers. Unlike previous literature, the forward-looking behavior of consumers is not driven solely by the anticipation of a marketing mix variable such as price or product quality, but instead by endogenous heterogeneity in consumer utility caused by current decisions and expectations about additional content. Our framework is applicable to a number of categories of experiential products, such as video games, TV shows, mobile applications, and educational products.

With our approach, we explain content consumption at the popular online game "World of Warcraft". We find three distinct consumer segments, with different preferences for current and future valuation of content: a relatively small segment of "serious" players that consume content at a fast pace, and two consumer segments that, although slower to consume content, still reach the same level of experience with the product as the former segment at a later stage in the product

time-effect dummy included in X_t.

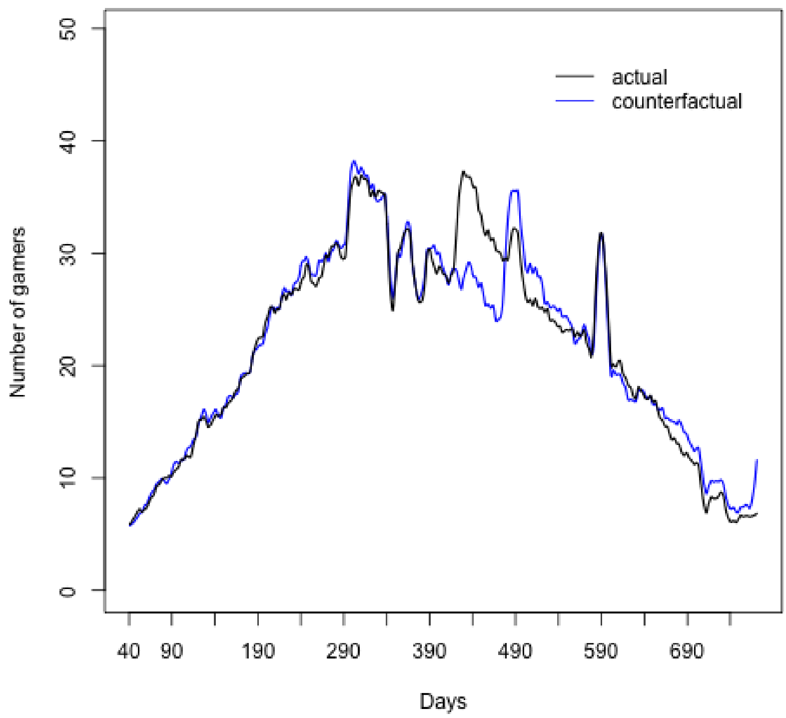

Figure 8: Evolution of the player participation in the actual scenario and in a scenario where the introduction of the fourth product update is postponed from day 421 to day 477.

lifecycle. This heterogeneity in speed of content consumption leads to interesting results when we tested alternative product update schedules. A schedule with later introduction of content leads to a positive impact on participation overall with less concentration of participation over time. With our approach, we quantify the value of a product update for the firm, measured in terms of consumer participation. This value can be translated into revenues based on price and subscription fees.

We believe that the social aspect of similar products can be further studied. In our application, we used a relative measure of the player involvement compared to the community. However, information about consumer decisions regarding belonging or not a group or community of players is now becoming increasingly available. Including these data in the estimation would allow us to quantify the contribution of "communities", instead of just the "individual", to the value of the product. We leave this question for future research.

References

[1] Arcidiacono, P. and J. B. Jones (2003), "Finite Mixture Distributions, Sequential Likelihood and the EM Algorithm," *Econometrica*, Econometric Society, vol. 71(3), pages 933-946, 05.

[2] Blizzard Entertainment (2008), "World of Warcraft subscriber base reaches 11.5 million worldwide", press release, available at http://us.blizzard.com/en-us/company/press/pressreleases.html?id=2847816.

[3] Chevalier, J. and A. Goolsbee (2009), "Are Durable Goods Consumers Forward Looking? Evidence from the College Textbook Market," *Quarterly Journal of Economics*, Vol. 124, November, 2009.

[4] Conlon, C. (2001), "A Dynamic Model of Costs and Margins in the LCD TV Industry," working paper.

[5] Danaher, P. and D. Mawhinney (2001), "Optimizing Television Program Schedules Using Choice Modeling," *Journal of Marketing Research*, 38, 3, 298- 312.

[6] Entertainment Software Association (2007), *Top 10 industry facts*, October 10th.

[7] Goettler, R. and R. Shachar (2001), "Spatial Competition in the Network Television Industry," *RAND Journal of Economics*, Winter 2001, 32(4), 624-656.

[8] Gordon, B. R. (2010), "A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry," *Marketing Science*, 28(5), 846–867.

[9] Hartmann, W. and V. Viard (2008), "Do Frequency Reward Programs Create Switching Costs? A Dynamic Structural Analysis of Demand in a Reward Program," *Quantitative Marketing and Economics*, 2008, vol. 6, issue 2, pages 109-137.

[10] Holbrook, M. B. and E. C. Hirschman (1982), "The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun," *Journal of Consumer Research*, Vol. 9, No. 2 (Sep.), pp. 132-140.

[11] Howard, J. A. and J. N. Sheth (1969), *The Theory of Buyer Behavior*, New York, John Wiley.

[12] Ishihara, M. (2011), "Dynamic Demand for New and Used Durable Goods without Physical Depreciation: The Case of Japanese Video Games," working paper.

[13] Misra, S. and H. Nair (2011), "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation," *Quantitative Marketing and Economics*, 9 (3), September, pg. 211-25.

[14] Nair, H. (2007), "Inter-temporal price discrimination with forward-looking consumers: Application to the US market for console video-games," *Quantitative Marketing and Economics,* Volume: 5, Issue: 3, Pages: 239-292.

[15] Norton, J. A. and F. M. Bass (1987), "A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High Technology Products," *Management Science*, September, 33, 9.

[16] Playlogic Entertainment Inc (2010), Form 10-Q, June 24th.

[17] Rao, A. (2011), "Online Content Pricing: Purchase and Rental Markets," Working Paper.

[18] Rideout, V., Roberts, D., & Foehr, U. (2005). Generation M: Media in the lives of 8–18 Year-olds. Washington, D.C.: Kaiser Family Foundation

[19] Song, I. and P. Chintagunta (2003), "A Micro-model of New Product Adoption with Heterogeneous and Forward Looking Consumers: Application to the Digital Camera Category," *Quantitative Marketing and Economics*, 1(4), pp. 371.

[20] Su, C.-L. and K. Judd (2011), "Constrained Optimization Approaches to Estimation of Structural Models," Working Paper, SSRN library.

[21] Williams, D., N. Yee, and S. Caplan (2008), "Who plays, how much, and why? Debunking the stereotypical gamer profile," *Journal of Computer-Mediated Communication*, 13, 993–1018, International Communication Association.

[22] Yang, S., V. Narayan, and H. Assael (2006), "Estimating the Interdependence of Television Program Viewership Between Spouses: A Bayesian Simultaneous Equation Model," *Marketing Science*, 25(4), 336-349.

[23] Yao, S., C. F. Mela, J. Chiang, and Y. Chen (2011), "Determining Consumers' Discount Rates with Field Studies," working paper, SSRN Library.