Customer Learning and Revenue-Maximizing Trial Design

Takeaki Sunada*

October 2, 2018

Abstract

Digital goods providers often offer free trials in order to familiarize customers with the product. I develop a structural model of customer learning-by-using to evaluate the profitability of two widely used trial configurations: limiting duration of free usage (i.e. "time-locked trial") and limiting access to certain features (i.e. "feature-limited trial"). Adopting a Bayesian learning framework, the model describes how product experience influences willingness to pay. It also allows identification of key factors behind learning and quantification of the trade-offs the firm faces in designing trials. I estimate the model using a novel data set of videogame users' play records. I find that in this setting, time-locked trials outperform feature-limited trials, and the revenue implication depends on the rate of demand depreciation during the trial period.

*Department of Economics, University of Pennsylvania (takeaki@sas.upenn.edu). I am extremely indebted to my committee, Aviv Nevo, Katja Seim and Eric Bradlow, for their guidance and support. The data set for this study is provided by the Wharton Customer Analytics Initiative (WCAI).

1 Introduction

Digital goods providers often offer trial versions of their products. For example, Spotify provides free trial periods of 60 days. Business and computational software such as Microsoft Office and Stata comes with a 30-day free trial. Videogame companies such as Electronic Arts and Sony offer several features from each game title for free. Reporting results from 305 software publishers, Skok (2015) finds that 62 percent of them generate some revenue from converting trial users to paid users, and 30 percent report that such conversion constitutes more than half of their revenue. A number of studies find that trial availability is positively associated with downloads of the paid version (Liu, Au and Choi 2014, Arora, Hofstede and Mahajan 2017). On the other hand, conversion rates vary significantly across firms and products. Spotify boasts that its conversion rate from the free trial to the paid version is around 27 percent, while Dropbox's conversion rate is around 4 percent; many other firms' conversion rates are only around 1 percent (Rekhi 2017). Brice (2009) and Turnbull (2013) both report that among the software providers they survey, the average conversion rate from a visit to the provider's website to purchase is *lower* with a trial than without it. Whether and to what extent a free trial boosts revenue therefore appears to depend on factors specific to each product and market.

In this paper, I develop an empirical framework to examine how different trial designs influence customers' adoption decisions and firm revenue. Designing the optimal trial is not a straightforward problem. A trial can be configured by limiting the duration of free usage ("time-locked trial") or by limiting access to certain features ("feature-limited trial").[1] If the trial is time-locked, the firm also needs to determine the duration of free usage. Similarly, when offering a feature-limited trial the firm needs to choose which features to include in the trial. In order to implement the optimal free trial, it is necessary to understand the trade-offs that each design entails.

I focus on one main factor through which the trial impacts revenue: customer learning-by-using (Cheng and Liu 2012, Wei and Nault 2013, Dey, Lahiri and Liu 2013). Digital goods are typical examples of experience goods. Advertising or other external information alone may not fully inform customers about their "match value": their individual valuation of the product.[2] When customers are risk averse, the existence of uncertainty lowers their willingness to pay for the product. A free trial thus informs customers about their true match value, increasing their willingness to pay. On the

[1] A "feature" refers to a general concept regarding "part of the product", such as game content, book chapter or news article. Since I consider a videogame, I use "content", "feature", and "game mode" interchangeably.

[2] Match values vary across customers due to different preference and needs, and ability to acquire product-specific skills.

other hand, it is essentially a free substitute for the full product, creating demand cannibalization. While providing a more generous trial product fosters better customer learning, it also increases the opportunity cost the trial imposes on the full product.

The costs and benefits of trial provision are associated with various aspects of customer learning. In this study, I consider four main factors: (1) the initial uncertainty around customer-product match value, (2) customer risk aversion, (3) speed of learning relative to demand depreciation, and (4) learning spill-overs across different features of the product. Initial uncertainty and risk aversion determine the magnitude of the drop in ex-ante willingness to pay. The speed of learning influences the profitability of a time-locked trial. As customers learn more quickly, shorter trial durations are necessary to facilitate learning, and demand depreciates less during the trial. Learning spill-overs across different features determine the effectiveness of a feature-limited trial. Since customers have free access to features included in the trial, willingness to pay consists only of customers' valuation of features excluded from the trial. Hence, in order for a feature-limited trial to increase willingness to pay, experience with trial features needs to be informative about excluded ones.

In order to evaluate the magnitude of each factor and predict how customers respond to different trial designs, I build and estimate a structural demand model for digital goods. Specifically, I embed the adoption decision of a durable product into a Bayesian learning framework. Customers make purchase decisions based on their willingness to pay, which is determined by their expected utility from future consumption stream (Ryan and Tucker 2012, Lee 2013). The expectation over future utility is conditional on customers' beliefs about their match value. Hence, both the magnitude of uncertainty reflected in the belief and risk aversion impact their willingness to pay. The model of Bayesian learning describes the process of how a customer updates her beliefs about her match value through product experience. In the model, each user maximizes her expected utility by choosing frequency of play, duration of each session and feature played in each session. Her uncertainty diminishes as she updates her belief. Learning spill-overs exist, in that an experience with one feature may help in updating the belief about her match value for other features. Moreover, a user may experiment with the product to explore her true match value; she takes into account future informational gains in choosing her actions. In order to capture the forward-looking nature of the decisions, I define the model as a dynamic programming problem (Erdem and Keane 1996, Che, Erdem and Öncü 2015). The solution of this problem provides a value function, which summarizes the customer's expected lifetime utility, determining her willingness to pay endogenously.

Customers' behavior under each trial design can be represented by combining the purchase decisions and the learning process in accordance with the trial design. In the case of no trial, the purchase decision precedes the entire learning process. Willingness to pay is determined by the value function evaluated with a prior belief. If a time-locked trial is provided, the learning model with the duration specified by the trial precedes the purchase decision. Once the trial expires, customers make a purchase decision based on their posterior belief. The posterior belief involves smaller uncertainty, increasing value from future sessions if customers are risk averse. However, initial free sessions no longer constitute willingness to pay, creating a trade-off for the firm. On the other hand, if a feature-limited trial is provided, trial users make a purchase decision after each trial session by comparing the value from switching to the full product and staying with the trial. The firm wants to include features in the trial that create large learning spill-overs, so that trial experience also reduces uncertainty about features excluded from the trial. Meanwhile, the firm also wants to exclude high-value features from the trial to prevent demand cannibalization. In either trial case, forward-looking customers take into account the option value of being able to make a better informed purchase decision in the future. Hence they may have a stronger incentive to experiment with the product during trial sessions.

I also account for other particularities of durable goods demand. First, customers may wait for future price drops (Stokey 1979). In order to capture this, the purchase decision is defined as an optimal stopping problem. Customers not only choose whether or not to purchase the product, but also determine the optimal timing of purchase (Nair 2007, Soysal and Krishnamurthi 2012). Second, in the model of learning, I account for other factors through which the past usage experience influences utility, such as novelty effect or boredom, and separately identify learning from them. Finally, I explicitly consider termination: permanent abandonment of the product. It determines the demand depreciation during the trial, influencing the trial's profitability.

I apply my framework to trial design of a major sports videogame. The videogame includes four features, which correspond to different content areas and are called "game modes". I use a novel data set of lifetime session records of 4,578 users. The records consist of the duration and the game mode selected at each session. The firm did not offer any free trial during the observation period, causing a sample selection problem: I only observe users having made a purchase without trial experience. I develop a procedure to estimate the population distribution of match values from an observation of such an endogenously truncated distribution. The estimation strategy also helps applying the current framework to other environments, where the available data set is typically subject to such sample selection problem.

I find that videogame users are risk averse, and their product valuation involves significant uncertainty. For example, consider a customer whose willingness to pay under her initial belief of her match value is \$50. The 95 percent confidence interval of her true willingness to pay is [\$21.80, \$87.90]. Users learn quickly. One additional session of a given game mode reduces the uncertainty about the match value with the mode by up to 63 percent. Meanwhile, learning spill-overs across different game modes are small. An additional session of a mode merely decreases the match value uncertainty of the other modes by 2 percent.

Given the estimated demand model, I evaluate the revenue implications of various trial designs. I find that in this setting, time-locked trials tend to increase revenue due to fast learning. In particular, providing five free sessions is the ideal design, increasing revenue by up to 2.5 percent. Meanwhile, I find that the revenue implications vary significantly with respect to the users' termination rate during the trial period. This implies that along with offering a free trial, the firm may want to incentivize users in order to decrease the trial termination rate. On the other hand, I find that any feature-limited trials without duration restrictions cannot increase revenue. This is due to a small number of modes in the product studied and small learning spill-overs; the opportunity cost of losing revenue from one mode always outweighs the gain from increased valuation for other modes. However, I also find that imposing extra feature restrictions on time-locked trials can boost revenue further by up to an extra 0.7 percentage point. This is because the extra feature restrictions widen the product differentiation between the full product and the trial; customers whose most preferred mode is excluded opt to make a purchase.

To my knowledge, this is the first empirical study that explores how the design of a free trial influences customers' adoption patterns. There exists an expansive theoretical literature on optimal trial provision when consumer learning exists (Lewis and Sapphington 1994, Chellappa and Shivendu 2005, Johnson and Myatt 2006, Bhargava and Chen 2012).[3] In particular, Dey, Lahiri and Liu (2013) and Niculescu and Wu (2014) study trade-offs the firm faces in providing time-locked and feature-limited trials, respectively. In both studies, the optimality condition depends on the aforementioned demand-side factors, calling for an empirical study to measure the magnitude of such factors. The methodology I propose does so by using typical data on customer engagement with the product. Hence, it can help firms design the optimal free trial.

Moreover, the model I develop considers customer learning in a durable goods environment,

[3] Trials may impact revenues through other channels. Cheng and Tang (2010) and Cheng, Li and Liu (2015) discuss network externalities that trial users create. Jing (2016) discusses the influence of trial provision on competition. Since neither factor is relevant to the product studied, I abstract away from these alternative stories.

offering a substantial departure from existing Bayesian models.[4] Unlike a perishable goods case, in this setting purchase and consumption are distinct, separable phenomena. For example, the product I study requires an outright purchase, and customers make only one purchase during the entire consumption stream. In such cases, the firm can manipulate *when* customers make a purchase over the course of learning; "timing of purchase" emerges as a new strategy variable of the firm. Moreover, forward-looking customers respond to different timing of purchase set by the firm by choosing *how to learn* during trial. For example, under a shorter free usage duration customers may have a stronger incentive for information acquisition and experiment more with the product. It endogenously determines the distribution of willingness to pay in response to the firm policy. My model is the first that explicitly considers such interactions beteen firm policy and customer actions specific to durable goods environments.[5]

This study also augments existing empirical studies concerning free trials of durable goods.[6] Similar to the optimal duration of a time-locked trial, Heiman and Muller (1996) study how the duration of product demonstration impacts subsequent adoption. Foubert and Gijsbrechts (2016) study implications of trial provision when the product involves quality issues. My study provides an empirical model that allows one to consider implications of different trial designs.

The study of optimal trial provision sheds a new light on the rapidly-growing "freemium" business model, where the firm offers part of its product for "free" and upsells "premium" components (Lee, Kumar and Gupta 2017). One of its main purposes is to facilitate learning from the free version and induce upsells: an objective similar to a feature-limited trial. Hence, understanding the mechanism of customer learning helps firms determine whether to adopt a freemium strategy.

This paper is structured as follows. In Section 2, I illustrate using a simple model how the mechanism behind customer learning affects firm revenue. In Section 3, I outline the data of videogame usage records. I also present supporting evidence for the existence of customer learning in the environment studied. In Section 4, I build the empirical model of digital goods adoption. I describe the identification and estimation strategy in Section 5. Estimation results and model fit

[4]There are numerous applications of Bayesian models to the repeat purchases of perishable goods, such as ketchup (Erdem, Keane and Sun 2008), yogurt (Ackerberg 2003), diapers (Che, Erdem and Öncü 2015) and detergent (Osborne 2011). Among others, physician learning about prescription drugs has a particularly large number of applications (Crawford and Shum 2005, Coscelli and Shum 2004, Chintagunta, Jiang and Jin 2009, Narayanan, Manchanda and Chintagunta 2005, Ching 2010, Dickstein 2018).

[5]Customers' endogenous usage adjustment in response to firm policy is studied in other contexts, such as multi-part tariff (Roberts and Urban 1988, Iyengar, Ansari and Gupta 2007, Narayanan, Chintagunta and Miravete 2007, Grubb and Osborne 2015, Goettler and Clay 2011). However, there customers merely respond to instantaneous price change and learning is affected only as a result; customers do not choose how to learn.

[6]More broadly, this study is associated with an empirical literature concerning how consumption experience in early stages influences future repeat behavior (Fourt and Woodlock 1960).

are discussed in Section 6. Using the estimated model I consider the optimal trial design in Section 7. Section 8 concludes and discusses possible future research areas.

2 An illustrative model of customer learning and firms' trade-offs

In this section, I introduce a simple model of customer learning and illustrate how each of the four factors outlined above impacts the optimal trial configuration. Consider a firm selling a videogame with two features, which provide flow utility v_1 and v_2, respectively. There is no complementarity across features and the utility from the full product is simply $v_1 + v_2$. For simplicity, I assume all customers have the same match value and receive the same utility. The product lasts for two periods. In the second period, the utility from both features decays by δ due to boredom. At the beginning of the first period, a customer faces uncertainty about her match value with the product. Her expected utility from each feature under uncertainty is given by $\mathbb{E}(v_i) = \alpha v_i$, for $i = \{1, 2\}$. $\alpha < 1$ is a parameter that captures the reduction of the utility due to the uncertainty in a reduced form way. α is low if a user faces large uncertainty or she is very risk averse. The uncertainty is resolved once she uses the product. When there is no free trial, the willingness to pay is equal to ex-ante expected utility from the whole product over two periods.[7]

$$\begin{aligned} U_N &= \mathbb{E}((v_1 + v_2) + \delta(v_1 + v_2)) \\ &= \alpha((v_1 + v_2) + \delta(v_1 + v_2)). \end{aligned}$$

Aside from not providing trial (N), the firm can either offer a time-locked trial (TL) or a feature-limited trial (FL). With TL, the customer uses the full product for free for one period and makes a purchase decision at the end of period 1. At the time of purchase, the customer learns her true match value but has only one active period remaining. Hence her willingness to pay, which equals the incremental utility from purchasing the full product, is

$$U_{TL} = \delta(v_1 + v_2).$$

With FL, the customer has free access to feature 1 and chooses whether to buy feature 2. Since she can only try feature 1, she may or may not be able to resolve the uncertainty about feature 2. I

[7] The customer knows that her uncertainty will be resolved in period 2. However, at the beginning she does not know the realization yet, and hence her period 2 utility is still in the expectation in the expression of U_N.

assume that learning occurs with probability γ. γ corresponds to the degree of learning spill-over. If two features are sufficiently similar, usage experience from one provides more information about the other and thus γ is high. I also assume that learning occurs well before period 1 ends. In this case, her willingness to pay is

$$U_{FL} = \begin{cases} v_2 + \delta v_2 & \text{with probability } \gamma, \\ \alpha(v_2 + \delta v_2) & \text{with probability } 1 - \gamma. \end{cases}$$

In order to maximize revenue from this customer, the firm first maximizes willingness to pay by choosing the scheme from N, TL and FL, and subsequently sets the price equal to it. When FL is provided, the price the firm sets is $p_{FL} = v_2 + \delta v_2$ and the customer purchases the full product only when learning occurs. Setting price $p_{FL} = \alpha(v_2 + \delta v_2)$ is dominated by choosing N and setting $p_N = U_N$. In Figure 1, I plot the area in which each of $\{U_{TL}, U_{FL}, U_N\}$ is the maximum of the three, and hence is the firm's optimal strategy. The result reflects the trade-offs described above.

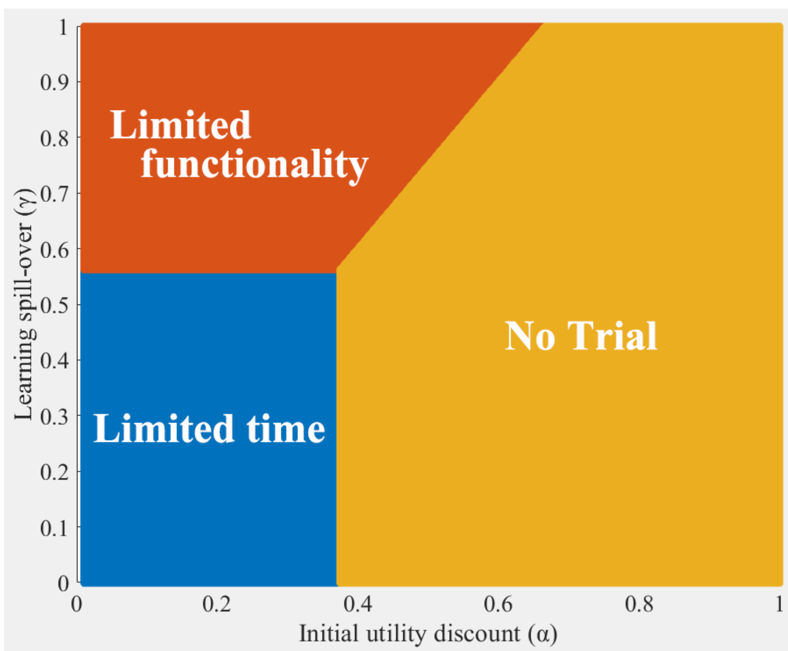

Figure 1: Optimality of trial schemes

Note: Each colored area represents the parameter range where each trial configuration is optimal. The figure is drawn assuming $v_1 = 1$, $v_2 = 2$ and $\delta = 0.6$.

With large α associated with small uncertainty and customer risk neutrality, providing trial is not optimal. When α is small, the optimal design depends on the relative size of α and γ. If learning spill-over γ is large, providing FL is optimal. On the other hand, if α is sufficiently small, it follows that the ratio $\frac{\alpha}{\delta}$ is also small; the second period utility remains high even after taking the boredom δ into account, so is the opportunity cost of providing free access in the second period. Providing TL is optimal in this case.

This model abstracts away many other factors at play, such as multi-period learning and customer heterogeneity. In addition, firms not only choose either TL or FL, but also the optimal free usage duration and features to include in the trial. In order to account for them, I develop a more realistic model in Section 4. Nonetheless, the factors discussed in this section remain the key drivers of the firm's trade-offs in the full model.

3 Empirical environment

3.1 Product studied and data description

In this study, I apply my framework to a major sports videogame. The game title is released annually with up-to-date real league data and enhanced graphics, and every version sells millions of units. The game operates on major gaming consoles such as Sony PlayStation series and Microsoft Xbox series. The game requires purchase of a game disk before it can be played. Hence customers make purchase only once. The game contains four features called "game modes". Each game mode correspond to a distinct content area, and users need to pick one mode whenever they play the game. In mode 1, users build a team by hiring players and coaches, and compete against rivals to win the championship. In mode 2, users simulate an individual sports player's career, in order to become the MVP. In mode 3, users pick a pre-defined team and play against other teams, skipping any team management. It is the most casual mode, involving simpler tasks and requiring less skills than other modes. Finally, mode 4 allows users to compete online and be ranked against other users. No predetermined play sequence exists and users can choose any modes from the beginning.

I observe a sample of 4,578 first-time users from the version released in 2014. The sample is randomly selected among U.S. based users who registered a user account during product activation. For each user, I observe the date of activation, which I assume to be the date of purchase, and the lifetime session records. Each session record consists of the time of start and finish, and the selected game mode. I augment my data with the game's weekly average market price, collected from a major price comparison website. The market price is the average of the prices listed on four major merchants: Amazon, Gamestop, Walmart and eBay. I assume that this market price is the purchase price.

A session, the unit of observation, is defined as a continuous play of one game mode. Upon turning the console on, the menu screen prompts users to select a game mode. Once they select one, a session starts. When they exit the mode and return to the menu screen, or shut down the console, the session ends. By definition of the session, each session consists of only one game mode.

The users in the data had no access to free trials; this is a set of users who made a purchase without trial experience.[8] As I show below, users in the data learn about the product by playing. Hence, no observation of free trial is necessary to identify learning. As I only observe a truncated subset of customers, I control for the sample selection problem during the estimation procedure.

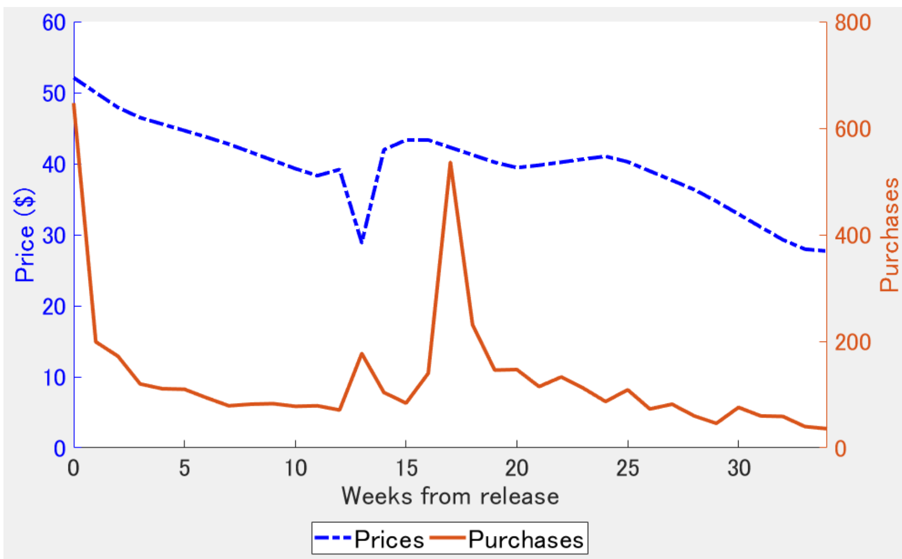

Figure 2: Prices and sales over time

Note: The prices are weekly average prices in the market. Sales are measured by the number of activations in the data.

In Figure 2, I show the history of sales and price over 35 weeks from the product release. Both follow a typical pattern of durable goods adoption: the highest price and the sales at the beginning, followed by a steady decline. In the 14th week, a lower price is offered due to Black Friday and a corresponding sales spike is observed. The 18th week is Christmas with a clear sales boost.

In Table 1, I present summary statistics of play records. Every mode is selected at a similar rate, indicating that these modes are horizontally differentiated. Each session lasts around an hour on average. Game mode 3 lasts shorter than the other modes, presumably because of its simplicity. Interval length between sessions is a measure of play frequency; shorter intervals indicates more

[8]Between 2012 and 2015 the firm changed its trial design every year, presumably in order to evaluate users' response. The no trial policy employed in 2014 is part of such experiment. Unfortunately, it is impossible to directly compare trial adopters and non-adopters using data from other years; the firm provided trials in a non-random manner, creating a sample selection problem that cannot be resolved with the current data.

Table 1: Summary statistics

		Mean	Std. Dev.
Mode 1	Choice probability	0.248	0.432
	Hours per session	1.200	1.200
Mode 2	Choice probability	0.316	0.465
	Hours per session	1.154	1.127
Mode 3	Choice probability	0.243	0.429
	Hours per session	0.570	0.818
Mode 4	Choice probability	0.193	0.394
	Hours per session	1.048	1.078
Session interval length (days)		2.743	4.358
Termination period (sessions)		30.742	43.795
Number of users		4,578	
Sample size (users × sessions)		145,317	

Note: Statistics are aggregated over all user-sessions. Choice probability is the aggregate proportion that each mode is selected. Its standard deviation is that of a user-session specific indicator variable, which is one if that mode is selected. Interval length between sessions is the number of days between two consecutive sessions. The termination period equals the number of total sessions each user played.

frequent play. On average, users play one session in every 2.7 days. The product life is relatively short. On average, users terminate after 31 sessions. Taken together, an average user remains active for slightly less than three months, plays 30 hours in total and then terminates.

In Figure 3, I show the heterogeneity of game mode selection across customers with different usage intensity, and its evolution over time. Each of the three bins represents users whose lifetime hours of play is in the bottom third (light users), middle third (intermediate users), and top third (heavy users) among those who remain active for at least 10 sessions. For each bin of users, each bar represents the proportion that each mode is selected in the initial 3 sessions, 4th-10th sessions, and sessions after that. Two empirical regularities are observed. First, the proportion varies across users with different intensities. For example, light users tend to play mode 3 more often than other users. This indicates that the distribution of match value over game modes may vary across users with different intensity.[9] Second, the proportion evolves over time. Mode 1 and 2 gain popularity, while mode 3 shrinks. This is indicative that the perceived match value evolves. These findings are consistent with customer learning, but also indicate the necessity to introduce

[9] I can interpret the shorter hours of play for mode 3 in two ways. There may exist ex-ante light users, who play short sessions and prefer mode 3. The other story is that mode 3 requires less time, and users who like mode 3 tend to play less hours. The two stories differ in the direction of causality. As long as light users receive lower utility and exhibit lower willingness to pay, I do not need to separate the two. In the data, users who prefer mode 3 tend to buy the game further away from the release at lower prices, indicating that they indeed have lower willingness to pay.

customer heterogeneity to account for the systematic difference across users.

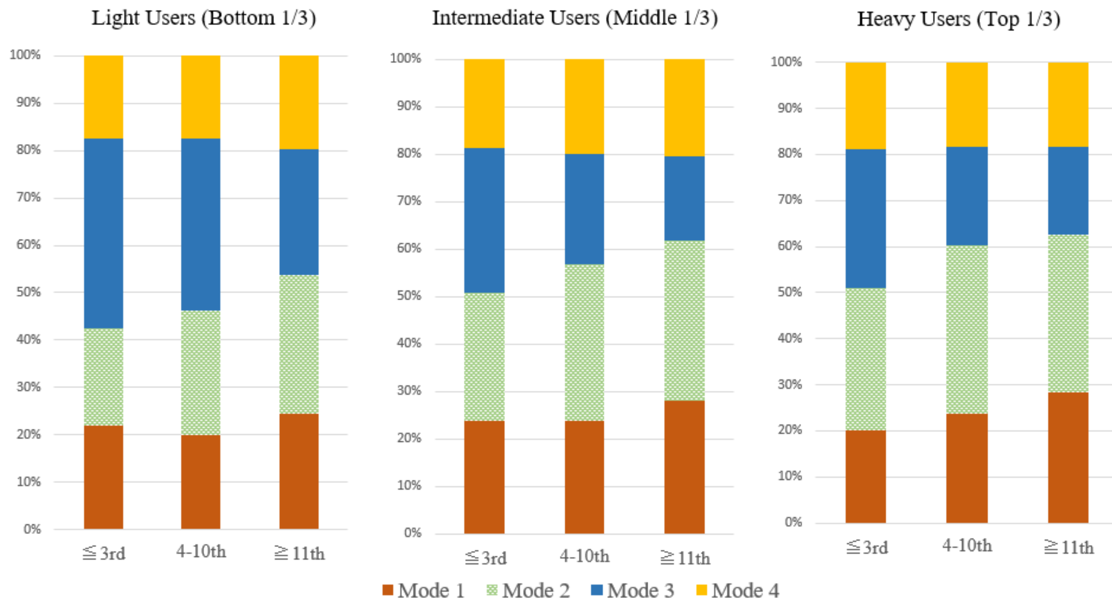

Figure 3: Evolution of game mode choice for each usage intensity

Note: Light, intermediate and heavy users are those whose lifetime hours of play are in the bottom, middle and top third of users. I exclude users who terminate within 10 sessions, in order to eliminate sample selection over time. For each bin of users, each bar represents the proportion that each mode is chosen in the first 3 sessions from the purchase, from 4th to 10th sessions, and 11th session and after.

In Figure 4, I present the evolution of session duration and session interval length for the same bins of usage intensity as in Figure 3. The usage pattern is nonstationary. On the one hand, the session durations increase over the first several sessions. Although this is consistent with learning story that the uncertainty resolution increases utility from play, there exist alternative explanations, such as skill acquisition and novelty effects. On the other, usage intensity declines in later sessions, likely due to boredom.[10] These findings imply that in order to correctly identify customer learning from the observed usage patterns, I need to control for other forms of state dependence. Users are heterogeneous both in session duration and play frequency; heavy users exhibit higher usage intensity and a slower decay than the others.

[10]Unlike role-playing games, sports games do not have pre-defined "ending" and users can keep playing it as long as they like.

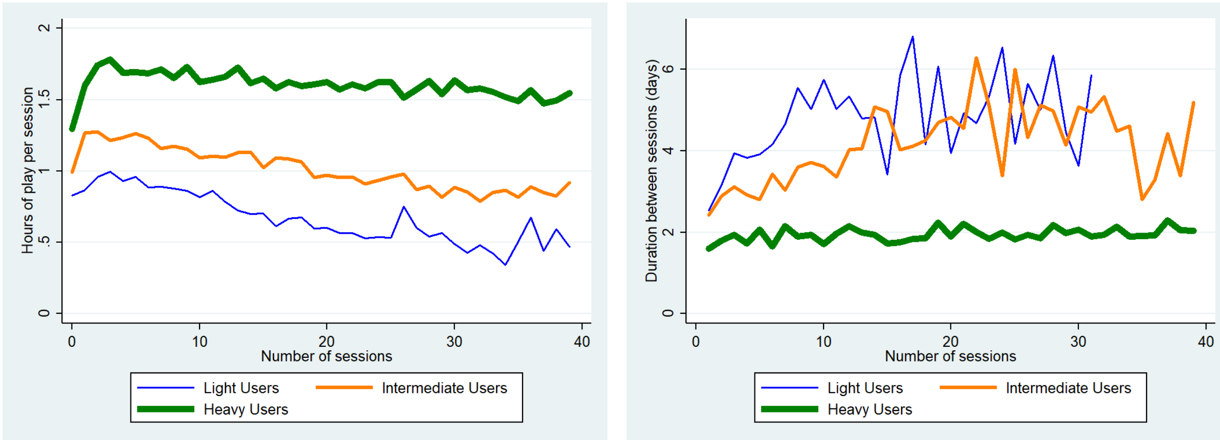

(a) Session duration

(b) Intervals between sessions

Figure 4: Evolution of session duration and intervals between sessions

Note: The session duration and interval length are the average across users for each usage intensity bin. The usage intensity is defined in the same way as in Figure 3. The figure is truncated at the 40th session.

3.2 Suggestive evidence of consumer learning

In this section, I discuss two data patterns that indicate the existence of customer learning-by-using. Other supportive evidence is discussed in the Appendix.[11]

High early termination rate On average, users terminate after 31 sessions. However, there exist many early dropouts. Figure 5 shows that 8.9 percent of users stop playing after the first session and 29.3 percent of users terminate within five sessions. Such a high early termination rate is also observed among heavy initial users: users whose first session duration is in top third. 6.6 percent of heavy initial users terminate after the first session, which is three times higher than their long-run average termination rate. Considering that most users purchase the game for around $40 to $50, some users are likely experiencing disappointment. Users who have high expectations about the match value may pay $40, only to realize their true match value is low and terminate early.[12]

Experimentation across game modes The existence of uncertainty implies that there is an option value from exploration; there is a possibility that the true match value with a mode is

[11] As shown by Chamberlain (1984), purely nonparametric identification between learning and customer heterogeneity is impossible. Hence, the validity of the argument that customer learning exists rests on how much the data pattern "intuitively makes sense" from the perspective of each story. Reassuringly, the observed patterns fit more naturally with customer learning. In the model, I impose restrictions on how heterogeneity can affect the evolution of utility, in order to identify learning. This assumption is often employed in the literature to disentangle state dependence from heterogeneity (Dubé, Hitsch and Rossi 2010, Ching, Erdem and Keane 2013).

[12] Heterogeneity in the speed of utility satiation may partly explain the patterns. Nevertheless, the fact that 9 percent of users play only one session after paying $40 strongly indicates the existence of uncertainty.

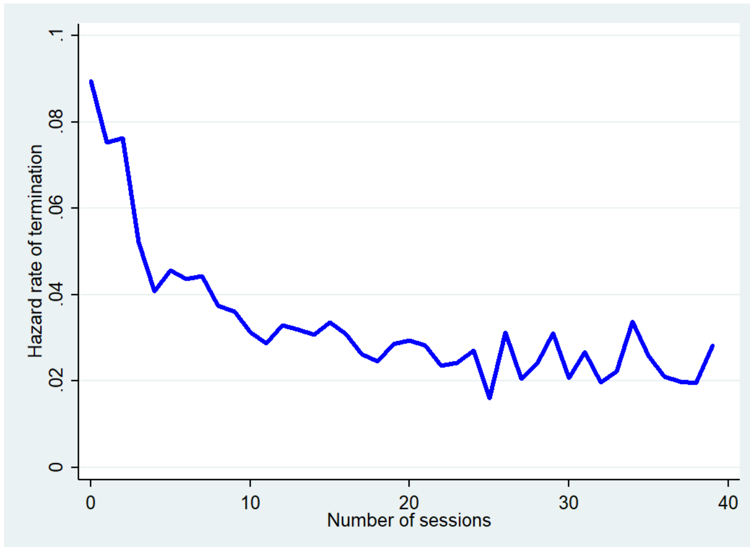

Figure 5: The evolution of the hazard rate of termination

Note: The hazard rate of termination at the t-th session is the proportion of users terminating after the t-th session among the active users. After 40 sessions, the hazard rate is stable.

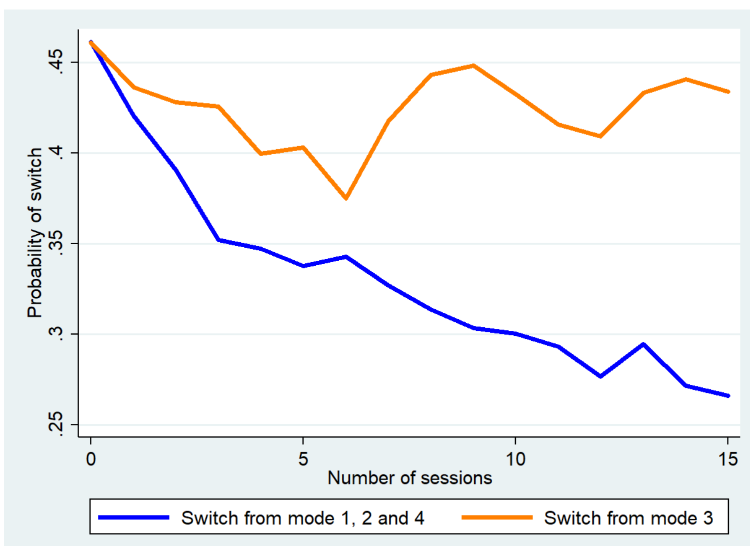

Figure 6: The evolution of the probability of switch

Note: The probability that users switch from mode m at the t-th session is calculated by the number of users who selected game mode $m' \neq m$ in the t+1-th session, divided by the number of users who selected mode m at the t-th session. Switches from modes 1, 2 and 4 follow very similar paths and I aggregate them for exposition.

quite high. This prompts users to experiment with each game mode, resulting in frequent switching across modes in the early stages of consumption. In Figure 6, I show the evolution of the probability that users switch modes after each session. Two different patterns emerge. The probability that users switch from modes 1, 2, and 4 to any other mode steadily declines as users accumulate more experience: a pattern consistent with experimentation. On the other hand, no experimentation of mode 3 seems to occur. These patterns are reasonable given that modes 3 is more casual and involves simpler in-game tasks. Indeed, as I show below, the estimated parameters support that match values with mode 3 involves smaller uncertainty.[13]

4 A structural demand model under no trial

In the previous sections, I showed that the usage patterns evolve in a way that is consistent with customer learning: popular game modes change over time; some users terminate very early; and users initially experiment across different modes. On the other hand, I found that usage experience may influence utility through other channels, such as boredom and novelty effects. Moreover, the existence of customer heterogeneity is strongly indicated.

In order to identify the mechanism behind learning and to predict customers' behavior under different counterfactual trial designs, I build and estimate a structural model of customers' adoption and learning-by-using. The structural model serves two purposes. First, it allows one to identify the learning mechanism, while controlling for other contaminating effects and customer heterogeneity. In particular, four factors that influence learning are explicitly modeled: magnitude of initial uncertainty, risk aversion, speed of learning, and learning spill-overs across different game modes. Second, the model is estimable using typical data on usage records of users *without free trial experience*. Hence, it provides implications for the trial profitability even without observing trial behavior.

In this section, I describe a model that corresponds to no trial case: the model I estimate with the data. Models under trial cases are discussed in Section 7. Under no trial, customers make a purchase decision prior to any learning-by-using. The initial adoption decision is based on their product valuation, which is equal to the sum of utilities they expect from future sessions, conditional on their prior belief about their match value. The utility from each subsequent session

[13] An alternative story is that customers merely have a taste for variety at the beginning. My standpoint is that learning and love of variety are not mutually exclusive, but that experimentation is a structural interpretation of variety seeking.

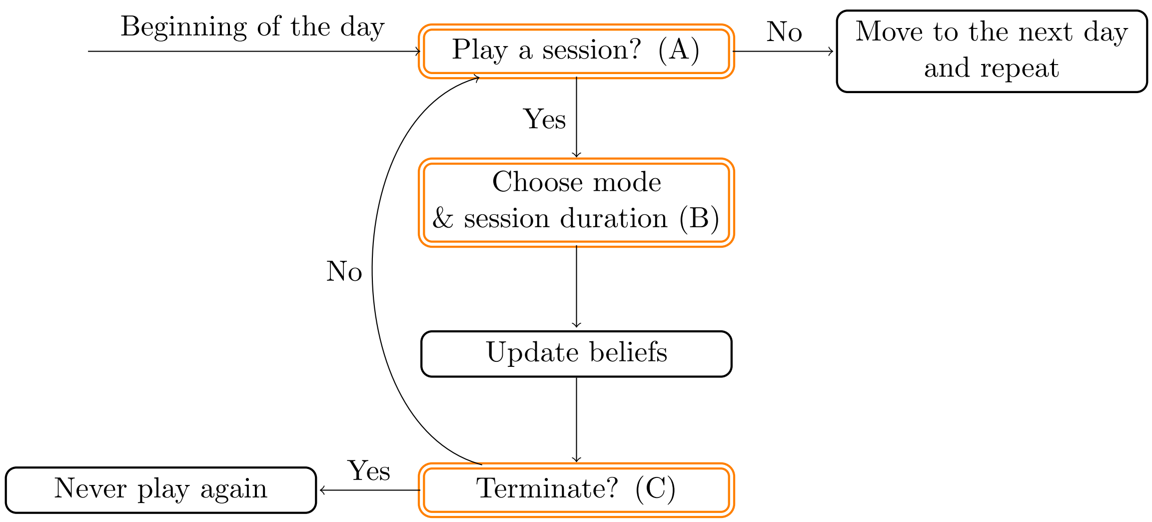

Figure 7: Timeline of the choices at each day

Note: Double-edged nodes A through C are decision nodes.

is endogenously determined through the decisions of product usage. Because of this structure, I first describe the model of learning-by-using and then the purchase decision.

4.1 The Bayesian learning model

The model of learning-by-using characterizes users' post-purchase play decisions. At each day, a user makes decisions according to a timeline described in Figure 7. The user first chooses whether to play a session. If she doesn't play one, she moves the the next day. If she plays one, she selects a game mode and chooses session duration. After a session, she receives a signal informative about her true match value from the selected mode, and updates her belief. At this point, the user may decide to permanently quit playing. I refer to this as termination. Termination is an absorbing state and she never makes any decisions again. Conditional on remaining active, the user chooses whether to play another session or move to the next day. She repeats this sequence until she terminates. In what follows, I first describe the user decisions during a session (Node B), and the decisions of play frequency and termination (Node A, C) afterward.

4.1.1 Selection of game modes and session duration (Node B)

At each session, users select a game mode and choose session duration. In order to allow for experimenting across modes, I assume that users are forward-looking and take into account future informational gains when selecting a game mode; users solve a discrete choice dynamic programming problem. Conditional on having selected a mode, users then choose duration of the session, which endogenously determines the flow utility from that mode.

Game mode selection A forward-looking user selects a game mode that maximizes the sum of her flow utility from playing and future informational return (Erdem and Keane 1996). In order to capture the nonstationary usage pattern presented in Figure 4, I assume that the problem has a finite horizon.[14] The optimal mode selection is summarized by the following value function.

$$V(\Omega_{it}) = \mathbb{E}[\max_{m_{it}} v(b_{it}, \nu_{imt}, h_t) + \mathbb{E}[\beta(\Omega_{i,t+1})V(\Omega_{i,t+1}) \mid \Omega_{it}, m_{it}] + \epsilon_{imt}\sigma_\epsilon],$$

where $\Omega_{it} = \{b_{it}, \{\nu_{imt}\}_{m=1}^{M}, h_t\}$; the state variables include b_{it}, i's belief about her match value at session t; ν_{imt}, the cumulative number of times that i chose mode m in the past t-1 sessions; and h_t, a weekend indicator, which is one for Saturday, Sunday and holidays. I denote the flow utility from the current session by $v(b_{it}, \nu_{imt}, h_t)$. The future informational gain is summarized by the continuation payoff $\mathbb{E}[\beta(\Omega_{i,t+1})V(\Omega_{i,t+1}) \mid \Omega_{it}, m_{it}]$. $\beta(\Omega_{i,t+1})$ is a discount factor between the current session and the next session. I discuss the definition of $\beta(\Omega_{i,t+1})$ below. The expectation of the continuation payoff is taken over an informative signal that the user receives after the current session. I assume that there exists a choice-specific idiosyncratic utility shock ϵ_{imt}, and that $\epsilon_{imt}\sigma_\epsilon$ follows type 1 extreme value distribution with variance σ_ϵ^2.[15] The choice probability of each mode hence follows the logit form.

$$P_m(\Omega_{it}) \quad = \quad \left(\frac{\exp\left(\frac{1}{\sigma_\epsilon}(v(b_{it}, \nu_{imt}, h_t) + \mathbb{E}[\beta(\Omega_{i,t+1})V(\Omega_{i,t+1}) \mid \Omega_{it}, m_{it}])\right)}{\sum_{m'} \exp\left(\frac{1}{\sigma_\epsilon}(v(b_{it}, \nu_{im't}, h_t) + \mathbb{E}[\beta(\Omega_{i,t+1})V(\Omega_{i,t+1}) \mid \Omega_{it}, m'_{it}])\right)} \right). \tag{1}$$

Experimentation occurs when a user chooses a mode that generates a lower flow utility than her current best alternative to gain a higher return in the future.[16] Since the problem is nonstationary, all the value functions and the optimal actions are a function of t in addition to the state Ω_{it}, which I suppress for notational simplicity.

[14] I assume that at $T = 100$ session all active users terminate. This is longer than the lifetime number of sessions of 93.27 percent of the users in the data.

[15] A commonly imposed normalization that $\sigma_\epsilon = 1$ is not necessary. The scale normalization is achieved by assuming that $f(b_{it})$ has no scaling coefficients. More details are provided in the Appendix.

[16] This trade-off is also found in the literature of Bandit models, and an index solution exists for this class of model with correlated arms (Dickstein 2018). Here I follow a dynamic programming approach, because the value function from the dynamic programming problem is a necessary input to the purchase decision.

Choice of session duration Now I derive $v(b_{it}, \nu_{imt}, h_t)$, the flow utility from mode m, as a result of optimal choice of session duration. I assume that having selected game mode m at session t, user i chooses duration of the session to maximize her expected utility specified as follows.

$$v(b_{it}, \nu_{imt}, h_t) = \max_{x_{imt}} f(b_{it}) x_{imt} - \frac{(c(\nu_{imt}) + x_{imt})^2}{2(1 + \alpha h_t)}. \quad (2)$$

x_{imt} is the session duration that user i chooses. I assume that the expected utility is a quadratic function of x_{imt}. f and c are functions that represent how marginal utility from playing an extra hour is affected by the belief b_{it} and the history of play ν_{imt}, respectively. Accumulation of usage experience influences utility through two channels. First, due to learning, users update their beliefs about match values and their utility evolves accordingly. This is captured through f. Second, aside from learning, usage experience may directly influence utility. c controls for such other factors. For example, any deterministic utility decay, such as satiation or boredom, implies that c is increasing in ν_{imt}. Likewise, due to novelty effects or skill acquisition, c may decrease in ν_{imt} for some range of t. Separating learning from other forms of state dependence that I discussed earlier corresponds to separately identifying c from learning parameters. Also, c partly captures the concept of demand depreciation. If the incremental utility from additional sessions decays quickly, providing initial free sessions incurs a large opportunity cost. Finally, users tend to spend more time in weekend, indicating that they may receive higher utility. This is captured by $\alpha > 0$.

The static expected utility maximization problem has a closed form solution, resulting in the following flow utility and the optimal session duration for each mode m.[17]

$$v(b_{it}, \nu_{imt}, h_t) = \frac{f(b_{it})^2 (1 + \alpha h_t)}{2} - f(b_{it}) c(\nu_{imt}), \quad (3)$$

$$x^*(b_{it}, \nu_{imt}, h_t) = f(b_{it})(1 + \alpha h_t) - c(\nu_{imt}). \quad (4)$$

Henceforth, I parametrize $f(b_{it})$ as follows.

$$f(b_{it}) = \mathbb{E}[\theta_{im}^{\rho} \mid \theta_{im} > 0, b_{it}], \quad (5)$$

where θ_{im} denotes a user-mode specific true match value. $f(b_{it})$ is specified as an expectation of

[17] I assume that users choose x_{imt} before playing. This ignores a possibility that learning from the current session influences the current session duration. However, this possible misspecification hardly impacts the estimate of willingness to pay. As the willingness to pay is defined by the sum of all future utility, the magnitude that learning shifts the utility from the current, single session is small relative to the entire valuation change from many future sessions.

θ_{im}^{ρ} conditional on the belief. $\rho > 0$ can be interpreted as the coefficient of risk aversion. $\rho < 1$ implies that utility is concave in the true match value θ_{im}, and taking expectation over θ_{im}^{ρ} results in utility diminution. On the other hand, c can be an arbitrary function such that $c(0) = 0$.

4.1.2 The decisions of play frequency and termination (Nodes A, C)

At decision nodes A and C, each user makes decisions of play frequency and termination. She compares value from playing to that from not playing at node A, and compares value from remaining active to that from terminating at node C. Instead of defining a full maximization problem, I take a reduced form approach to model them. Specifically, I impose the following two assumptions; (1) users' decisions are based only on the state Ω_{it} at any decision nodes located between sessions t-1 and t, and (2) decisions are influenced by an idiosyncratic shock, such that the optimal decision is representable by a probability distribution over each of the available alternatives. This encompasses many specifications of decision rules that involve an idiosyncratic utility shock, some of which I discuss in the Appendix. I denote the probability that user i plays her t-th session on a given day by $\lambda(\Omega_{it})$, and the probability that user i remains active after session t by $\delta(\Omega_{i,t+1})$. I treat these probability distributions as model primitives.

Given the structure of the decisions of frequency and termination, I derive the formula for $\beta(\Omega_{i,t+1})$: the discount factor between session t and t+1. Assuming that users discount future utility by β per one day, $\beta(\Omega_{i,t+1})$ is obtained as the expected discount factor between the date that session t is played and the date that session t+1 is played. The expectation is over whether the user remains active after session t, and when she plays session t+1; because the optimal action at each decision node depends on an idiosyncratic shock that only realizes at that node, a user's future actions are stochastic to herself. Formally, $\beta(\Omega_{i,t+1})$ is characterized as follows.

$$\begin{aligned}\beta(\Omega_{i,t+1}) &= \delta\lambda + \delta(1-\lambda)\lambda\beta + \delta(1-\lambda)^2\lambda\beta^2 + ... \\ &= \delta(\Omega_{i,t+1})\frac{\lambda(\Omega_{i,t+1})}{1-(1-\lambda(\Omega_{i,t+1}))\beta}. \end{aligned} \tag{6}$$

The intuition is as follows. After session t the user remains active with probability $\delta(\Omega_{i,t+1})$. If staying active, she plays session t+1 on the same day with probability $\lambda(\Omega_{i,t+1})$, on the next day with probability $(1-\lambda(\Omega_{i,t+1}))\lambda(\Omega_{i,t+1})$, incurring the daily discount factor β, and so on.

4.1.3 State variables and their evolution

Match value and its learning-by-using I denote the true match value of customer i with the game by a vector $\theta_i = \{\theta_{i1}, \theta_{i2}, ..., \theta_{iM}\}$, where M is the number of modes available in the full product. Users are heterogeneous in their match values. I assume that θ_i follows multivariate normal distribution; $\theta_i \sim N(\mu, \Sigma)$, where $\mu = \{\mu_1, \mu_2, ..., \mu_M\}$ is the average match value of the population and Σ is an arbitrary variance-covariance matrix. Correlations between match values across different modes can be either positive or negative. Heavy users play all modes more extensively than light users, generating positive correlations. On the other hand, users who like mode m tend to play only mode m and not other modes. This generates negative correlations.

Upon arrival at the market, the customer does not know the realization of θ_i, and has a rational expectation about it; her prior about the distribution of θ_i is equal to the population distribution of match values.[18] In addition, she receives a vector of initial signals $\tilde{\theta}_{i0} = \{\tilde{\theta}_{i10}, \tilde{\theta}_{i20}, ..., \tilde{\theta}_{iM0}\}$. The initial signal represents any information or impression that a user has about the product ex-ante. It creates heterogeneity in the initial perceived match value. I assume that the initial signal is independent across game modes, and is normally distributed conditional on the user's true match value; $\tilde{\theta}_{im0} \mid \theta_{im} \sim N(\theta_{im}, \tilde{\sigma}_m^2)$. Henceforth, I denote the diagonal matrix of the variance of the initial signal by $\tilde{\Sigma}$. The customer forms an initial belief as a weighted average of the prior distribution and the received signal in a Bayesian manner.

$$\begin{aligned} &\theta_i \mid \tilde{\theta}_{i0} \sim N(\mu_{i1}, \Sigma_1), \\ &\text{where } \mu_{i1} = \mu + \Sigma(\Sigma + \tilde{\Sigma})^{-1}(\tilde{\theta}_{i0} - \mu), \\ &\qquad \Sigma_1 = \Sigma - \Sigma(\Sigma + \tilde{\Sigma})^{-1}\Sigma. \end{aligned} \tag{7}$$

The initial belief is thus represented by $b_{i1} = \{\mu_{i1}, \Sigma_1\}$.

When a user plays game mode m at each session t, she receives a signal informative about her true match value for that mode. I assume that the signal s_{imt} is normally distributed around the true match value.

$$s_{imt} \mid \theta_{im} \sim N(\theta_{im}, \sigma_s^2).$$

[18] This assumption does not allow for possible bias in the initial belief. The bias in the belief, if it exists, is not separately identified from other forms of deterministic utility evolution, and hence is currently subsumed in $c(\nu_{imt})$.

I assume that the variance of the signal σ_s^2 remains the same over time. Introducing a time-varying signal distribution makes the model computationally intensive, and hence I opt to maintain a simple structure. Given the realized signal, the user updates the belief following the Bayesian formula.

$$\mu_{i,t+1} = \mu_{it} + \Sigma_{it} Z'_{it} (Z_{it} \Sigma_{it} Z'_{it} + \sigma_s^2)^{-1} (s_{imt} - \mu_{imt}), \tag{8}$$

$$\Sigma_{i,t+1} = \Sigma_{it} - \Sigma_{it} Z'_{it} (Z_{it} \Sigma_{it} Z'_{it} + \sigma_s^2)^{-1} Z_{it} \Sigma_{it}, \tag{9}$$

where Z_{it} is a 1 by M vector whose m-th element is one and zero elsewhere.[19]

This learning structure captures all four factors that determine firm-side trade-offs described earlier. The magnitude of initial uncertainty is determined by the variance of the initial belief Σ_1. Customers know that their beliefs involve an error and hence they face a risk of mismatch. When customers are risk averse, $\rho < 1$ in Equation (5) and the expected utility from future sessions is lowered. On the other hand, the speed of learning is captured by σ_s. When σ_s is small, signals is more precise and uncertainty diminishes more quickly. Finally, learning spill-overs are determined by match value correlation Σ. If the match values for two game modes are highly correlated, a signal received from one mode also helps update the belief for the other. These are the key parameters that determine the profitability of each trial design in the counterfactual exercise. Willingness to pay is likely to increase due to learning when customers are sufficiently risk averse, magnitude of initial uncertainty is large and learning is quick relative to the utility depreciation.

Evolution of other state variables ν_{imt} evolves deterministically; $\nu_{im1} = 0$ for all m, and $\nu_{im,t+1} = \nu_{imt} + 1$ if m is chosen at session t, and $\nu_{im,t+1} = \nu_{imt}$ otherwise. The weekend indicator is i.i.d, and it is 1 with probability 2/7 and zero with probability 5/7. This stochastic weekend arrival helps reduce the dimension of state variables; deterministic weekend arrival requires to keep track of the day of the week in the state.[20] This completes the description of the model for usage. This problem is solvable by backward induction. The solution consists of the optimal decision rule and the associated value function at each state Ω_{it}.

[19] This Bayesian updating structure of a normal distribution does not require to keep track of $\Sigma_{i,t}$ in the state space. Instead, it suffices to keep the mean belief μ_{imt} and the number of times each option is taken in the past ν_{imt} (Erdem and Keane 1996). This allows one to reduce the effective state space to $\Omega_{it} = \{\{\mu_{imt}, \nu_{imt}\}_{m=1}^{M}, h_t\}$.

[20] The derivation of $\beta(\Omega_{i,t+1})$ in Equation (6) ignores the fact that h_t evolves day by day even without playing a session. In practice, I replace relevant $\lambda(\Omega_{i,t+1})$ in Equation (6) with its expectation over the realization of h_{t+1}.

4.2 Purchase decisions

When there is no free trial, each customer makes purchase decisions without any playing experience. A user's product valuation is represented by her ex-ante value function $V(\Omega_{i1})$: the sum of the utility she expects from the product in the future, evaluated at the initial state Ω_{i1}.

The purchase decisions proceed as follows. I assume that the market consists of N customers. They are heterogeneous, in that $V(\Omega_{i1})$ is customer-specific. In line with the frequency of the price data, I assume that one period in the adoption model is one week. At each week τ, a fraction λ_τ^a of the customers randomly arrive. In light with the fact that the price of the product steadily declines over time, I allow customers to delay their purchase to buy at lower prices (Nair 2007). Each customer makes a purchase decision by comparing the value from buying to that from waiting for a price drop. If she makes a purchase, she quits the market and starts using the product in a way described above. If she does not make a purchase, she comes back to the market in the following week and makes the decision again. I assume that the product is available for 52 weeks after the release date, and hence waiting beyond 52 periods generates zero payoff. A new version of the game is released at week 52, when the sales of older version essentially end.

The value function associated with the optimal stopping problem at week τ is

$$V_{ip}(\Omega_{i1}, p_\tau) = \mathbb{E}[\max\{V(\Omega_{i1}) - \eta_i p_\tau + \epsilon_{1i\tau}\sigma_p,\ \beta V_{ip}(\Omega_{i1}, p_{\tau+1}) + \epsilon_{0i\tau}\sigma_p\}],$$

where p_τ is the current price. If the customer buys the product, she receives the value $V(\Omega_{i1})$ and pays p_τ. If she does not buy at week τ, she receives a continuation payoff of staying in the market $V_{ip}(\Omega_{i1}, p_{\tau+1})$. I assume perfect foresight for the future prices.[21] $\epsilon_{i\tau}\sigma_p$ is i.i.d, and follows type 1 extreme value distribution with variance σ_p^2.[22] I do not model social learning, and hence the value from purchase $V(\Omega_{i1})$ remains constant over time. Hence, the incentive to delay the purchase solely comes from lower prices in the future. I assume that η_i follows log-normal distribution with mean μ_η and variance σ_η^2. For simplicity, I assume that η_i is independent from θ_i. The probability that customer i makes a purchase at week τ follows the logit form.

$$P_{ip}(\Omega_{i1}, p_\tau) = \frac{\exp\left(\frac{1}{\sigma_p}(V(\Omega_{i1}) - \eta_i p_\tau)\right)}{\exp\left(\frac{1}{\sigma_p}(V(\Omega_{i1}) - \eta_i p_\tau)\right) + \exp\left(\frac{\beta}{\sigma_p} V_{ip}(\Omega_{i1}, p_{\tau+1})\right)}.$$

[21] Assuming instead that customers form expectations based on the prices of other game titles hardly changes the result. As in Figure 2, the prices follow a quite typical path and it is easy for customers to forecast price patterns.

[22] Since $V(\Omega_{i1})$ is already scale-normalized, I do not need to normalize σ_p^2.

The model is solvable by backward induction. The customer's willingness to pay for the product under no trial is defined by $\frac{V(\Omega_{i1})}{\eta_i}$: the value of the product measured in dollars.

4.3 Multiple segments

In addition to heterogeneities with respect to the true match value θ_i, the belief b_{it} and the price coefficient η_i described earlier, I allow for the existence of multiple segments $r = \{1, 2, ..., R\}$ with different population-level parameters. In particular, I allow the vector of mean match value μ and the variance of utility shock in the mode choice σ_ϵ to be heterogeneous. I denote segment-specific parameters with subscript r. I also let the variance of the initial belief be heterogeneous, denoted by $\Sigma_{1r} = \kappa_r(\Sigma - \Sigma(\Sigma + \tilde{\Sigma})^{-1}\Sigma)$ with $\kappa_1 = 1$. Introducing multiple segments allows for more flexible representation of customer heterogeneity; heterogeneity in μ allows for the existence of ex-ante heavy and light user segments, and heterogeneity in σ_ϵ and Σ_1 adds flexibility in fitting game mode selection of users with different usage intensity. The probability that each user belongs to segment r is denoted by ξ_r. This completes the model representation under no trial case.

5 Identification and estimation

5.1 Parameter identification

In this section, I outline the intuition behind identification of the key parameters. In particular, I address two main challenges: (1) separating learning from other forms of state dependence, such as boredom, and (2) separate identification of each of the four factors of learning. A formal identification argument and identification of other parameters are presented in the Appendix. For simplicity, here I consider a case with a single segment; $R = 1$. I denote the states observable to a researcher by $\bar{\Omega}_{it} = \{\{\nu_{imt}\}_{m=1}^{M}, h_t\}$. The difference from Ω_{it} is that $\bar{\Omega}_{it}$ does not include the belief b_{it}, which is unobservable to a researcher.

I first separately identify learning and other factors using observation of $x^*_{imt}(\bar{\Omega}_{it})$: the session duration at each state. "Identification of learning" refers to the identification of the distribution of mean beliefs μ_{imt} for each m at each $\bar{\Omega}_{it}$. Once it is identified, learning — how beliefs evolve across states — is identified immediately. On the other hand, other forms of state dependence are captured by $c(\nu_{imt})$. Separate identification of the two relies on two features of the model. First, $c(\nu_{imt})$ is a deterministic function of usage history ν_{imt}. Hence, users who share the same history ν_{imt} face the same $c(\nu_{imt})$. Second, the evolution of μ_{imt} due to learning is stochastic, involving initial signal $\tilde{\theta}_{i0}$

and post-session signals s_{imt}. Moreover, because of rational expectations, users' mean beliefs stay the same on average after receiving an incremental signal: $\mathbb{E}(\mu_{im,t+1}|\mu_{imt}) = \mu_{imt}$. These conditions imply that conditional on the history of usage up to session t, how the *average* session duration x^*_{imt} evolves from state $\bar{\Omega}_{it}$ to $\bar{\Omega}_{i,t+1}$ is solely attributed to the evolution of $c(\nu_{imt})$; evolution of μ_{imt} due to learning cannot influence the average behavior because of rational expectation. On the other hand, how the *variance* of x^*_{imt} across users evolves is solely attributed to the evolution of μ_{imt}; $c(\nu_{imt})$ cannot influence variance because users sharing the same history has the same $c(\nu_{imt})$. This achieves separate identification of the distribution of μ_{imt} due to learning and $c(\nu_{imt})$. In the Appendix, I provide a formal argument of the identification of the distribution of μ_{imt} at each $\bar{\Omega}_{it}$.[23]

Next I consider separate identification of the four factors of learning. Each of the four factors is captured by parameters $\{\rho, \tilde{\Sigma}, \sigma_s^2, \Sigma\}$: risk aversion by ρ, the magnitude of initial uncertainty by $\tilde{\Sigma}$, speed of learning by σ_s^2 and learning spill-overs by Σ, respectively. The distribution of μ_{it} at each $\bar{\Omega}_{it}$, which I identified in the previous paragraph, is sufficient to identify Σ, $\tilde{\Sigma}$ and σ_s^2. Specifically, I use the evolution of $Var(\mu_{it} \mid \bar{\Omega}_{it})$. The intuition is as follows. In the initial session, each user forms the initial belief using her prior and the initial signal; the distribution of beliefs μ_{i1} reflects Σ and $\tilde{\Sigma}$. As she learns her true match value θ_i, the distribution of μ_{it} converges to that of θ_i, which reflects only Σ. Moreover, the speed of convergence is determined by the precision of the signal σ_s^2. Hence, by observing the variance of the belief at the early stage of consumption, that in the long-run, and the speed of convergence, one can identify Σ, $\tilde{\Sigma}$ and σ_s^2.

Finally, the identification of ρ relies on the intertemporal switching across game modes during the initial experimenting periods. Since all the other learning parameters and other channels $c(\nu_{imt})$ are identified solely from the observation of session durations, the only remaining parameter to fit the initial switching patterns is ρ. Intuitively, given the belief users experiment with smaller number of modes if ρ is small. When ρ is small, customers are more risk averse and hence trying a new, unfamiliar game mode is more costly.

5.2 Estimation

I estimate the model using simulated method of moments. Given a set of candidate parameters, I first solve the model by backward induction. In order to account for continuous state space, I use the discretization and interpolation scheme (Keane and Wolpin 1994). I then simulate sequences of

[23] This argument assumes away sample truncation due to switching and terminating. In the Appendix I show how to accommodate these factors.

actions according to the optimal policy predicted by the model; I draw a set of true match values, initial signals, and post-session signals and record the predicted actions. The set of simulated users serves as pseudo-data. The parameters are estimated such that the pseudo-data obtained this way match most closely with the real data, according to pre-selected moments. Formally, for a vector of parameters θ the estimator $\hat{\theta}$ is obtained by the following minimization problem.

$$\hat{\theta} = \arg\min_{\theta} m_k(\theta)' \hat{V}^{-1} m_k(\theta),$$

where $m_k(\theta)$ is a vector, with rows containing the difference between the data and model moments. $\hat{V}$ is a weighting matrix.[24]

The set of moments is selected to closely follow my identification strategy. For the model of usage, at each observed history of play $\{\nu_{imt}\}_{m=1}^{M}$, I take as moments (1) the probability that each game mode is selected, (2) the probability that a user switches modes from the previous session, (3) the mean and variance of the duration of each session, (4) the average interval length until the next session and (5) the probability of termination. Since the number of possible paths grows as t becomes larger, there are only 172 states that I have a sufficient number of samples to satisfactorily compute these statistics. Most of them are located at the early stages of usage history.[25] In order to augment the set of moments in later periods, I calculate the same statistics at each session t, but aggregated across different histories of game mode selections $\{\nu_{imt}\}_{m=1}^{M}$ and use them as moments. Also, in order to exploit the variation of usage patterns across users with different usage intensity, I calculate the above moments conditional on multiple bins of usage intensity. I describe the construction of the bins in the Appendix. Finally, I add as an extra set of moments the difference of the average session duration between weekdays and weekends, and the probability that users play multiple sessions within a single day. These extra moments are designed to aid the identification of α and λ, respectively.

The data used to identify the model of adoption are the rate of adoption at each week τ from the release until the 16th week, which is two weeks before Christmas. The empirical rate of adoption is equal to the proportion of customers making a purchase at each week in the data, multiplied by the market share of the product. Market share of this product is 28.1 percent, which is calibrated using an external data set described in the Appendix. I do not use the rate of adoption on and

[24] As a weighting matrix, I use a diagonal matrix, whose {k,k} element corresponds to the inverse of the mean of the sample moment. This works to equalize the scale of the moments by inflating the moments that have a smaller scale (e.g. choice probability) while suppressing the moments with a larger scale (e.g. interval length between sessions).

[25] I use moments from the states with more than 30 observations.

after Christmas. This is because users activating the product on Christmas may have received it as a gift, and hence including their activation as a purchase would possibly bias the estimate of the price coefficient. In total, I have 7,375 moment conditions to estimate 47 parameters.

In order to consistently estimate the model parameters, I need to control for the sample selection problem; I only observe users who purchased the game without trial experience. In the context of simulated method of moments, the sample selection implies that in order to calculate moments regarding product usage, I need to construct a pseudo-data set of users *who made a purchase*: a data set comparable to the real data. I achieve this through the following procedure. At each parameter value, I first draw a set of potential customers from the population distribution, each with willingness to pay $V(\Omega_{i1})/\eta_i$. Using the adoption model, I calculate the probability that each simulated customer makes a purchase. Customers with high willingness to pay receives a high probability, and vice versa. I then simulate each customers' adoption decision according to that probability. I use the subset of customers predicted to make a purchase as the pseudo-data to calculate moments of the usage model. As I move through iteration rounds, both the set of users included in this subset and their actions are updated simultaneously, until the usage patterns of the selected subsample perfectly match with the data. The step-by-step description of my estimation procedure and construction of all the moments are detailed in the Appendix.

$c(\nu_{imt})$, $\lambda(\Omega_{it})$ and $\delta(\Omega_{it})$ are specified as a quadratic function with respect to the number of past sessions, whose coefficients can vary across users with different match values. I assume that the variance of the initial signal is proportional to the variance of the true type; $\tilde{\sigma}_m^2 = \kappa\sigma_m^2$. Also, without loss of generality I normalize the average match value of segment 1 customers with game mode 3 to 30, and define other parameters relative to it. Cocerning the number of discrete segments, I assume $R = 2$. The customer arrival process λ_τ^p is specified as a uniform arrival rate λ_u^a and the initial mass of arrival at the release date λ_0^a. I assume that the timing of arrival is independent from the location of initial beliefs. Assuming that N potential customers exist in the market, I can normalize $\lambda_0^a = 1$ and estimate only λ_u^a as the rate of arrival in the later weeks relative to the initial week.[26] Market size N is calibrated outside the model and it is equal to the installed base of consoles, multiplied by the share of sports games among all the game sales.

[26]While these assumptions are not very flexible, the data provide only 16 points of observation, from which I identify both the arrival process and the distribution of price coefficient. Hence, I opt for a simple process to avoid identification issues.

6 Estimation Results

In order to conduct model validation exercises, I randomly split the 4,578 users in the data into an estimation sample of 3,778 users and a holdout sample of 800 users. In this section, I first present the estimated parameters, and the model fit with all 4,578 users in the data. The model fit with the holdout sample is provided in the Appendix. I then use the estimated model to show how mechanism behind learning is at play in the current setting.

6.1 Parameter estimates of usage model

Table 2: Parameter estimates

Parameters		Estimates	Std.error
Risk aversion ρ		0.215	0.034
Holiday effect α		1.396	0.023
Distribution of match value			
Std.errors	σ_1	69.149	0.665
	σ_2	83.374	2.464
	σ_3	35.881	3.513
	σ_4	64.096	2.145
Correlations	ρ_{12}	0.510	0.059
	ρ_{13}	0.262	0.093
	ρ_{14}	0.588	0.035
	ρ_{23}	0.207	0.141
	ρ_{24}	0.515	0.050
	ρ_{34}	0.541	0.069
Initial signal var κ		0.323	0.095
Post-session signal s.e σ_s		28.497	0.737

(a) Common parameters

Parameters		Segment 1			Segment 2	
		Estimate	Std.error		Estimates	Std.error
Mean match value	μ_{11}	24.062	16.359	μ_{12}	107.106	3.604
	μ_{21}	33.844	8.090	μ_{22}	95.780	3.593
	μ_{31}	30	0	μ_{32}	103.449	1.945
	μ_{41}	32.089	6.616	μ_{42}	97.070	6.810
Initial uncertainty	k_1	1	0	k_2	7.758	1.415
Logit shock s.e	$\sigma_{\epsilon 1}$	2326.129	37.541	$\sigma_{\epsilon 2}$	3227.733	214.038
Proportion of seg 1	ξ_1	0.578	0.016			

(b) Segment-specific parameters

Note: μ_{31} and k_1 are normalized. Standard error is calculated by 1,000 bootstrap simulations.

In Table 2, I present selected parameter estimates of the usage model. The standard errors are simulated using 1,000 sets of bootstrapped data set, each of which is obtained by randomly re-sampling users from the original data with replacement. In Table 2a, I show the estimates of the parameters common across all users. The coefficient of risk aversion is $0.215 < 1$, indicating significant risk aversion. The standard error of the match value distribution is much smaller for game mode 3; because of the simplicity of the mode, match values vary little across users. The estimate is consistent with no experimentation of mode 3 in Figure 6. Also, correlation coefficients are all positive. A high match value for one mode implies a high match value for another. However, as I show below, the magnitude of the correlation is not high enough to generate much learning spill-over. In Table 2b, I present segment-specific parameters. Each of the two discrete segments respectively captures the behavior of light users and heavy users. All the parameters for segment 2 are inflated to capture the large gap of usage intensity between light and heavy users.[27]

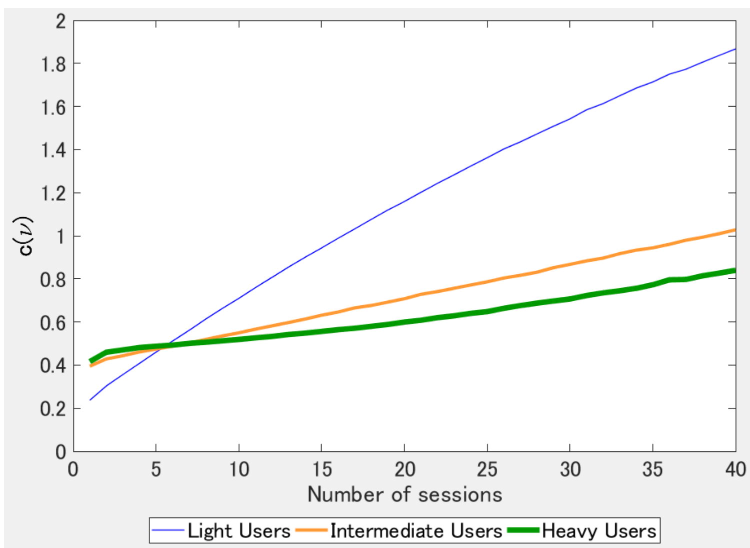

Figure 8: Evolution of estimated $c(\nu_{imt})$

Note: Each line corresponds to the average evolution of $c(\nu_{imt})$ of users who belong to each bin of usage intensity. They are calculated using 50,000 simulated sequences of actions. Bins of usage intensity are defined as in Figure 3.

In Figure 8, I present the evolution of utility due to other forms of state dependence captured in $c(\nu_{imt})$ across users with different usage intensity. Higher $c(\nu_{imt})$ implies lower marginal utility from an extra hour of play. Utility monotonically decays over time; the increase of utility due

[27] While the estimated magnitude of initial uncertainty that segment 2 faces is disproportionately high, this merely reflects the tight curvature of the utility due to small ρ. Since the flow utility is quite flat at a high match value, in order to capture the fact that the uncertainty also reduces heavy users' initial utility, the variance of the belief needs to be magnified accordingly.

to skill acquisition or novelty effects does not seem to exist in this setting. The utility of heavy users tend to decay slower than others, consistent with Figure 4. As I detail in the Appendix, all the parameters of $c(\nu_{imt})$ are precisely estimated and significantly different from zero; aside from learning, other forms of state dependence exists.

6.2 Model fit and implications for usage pattern

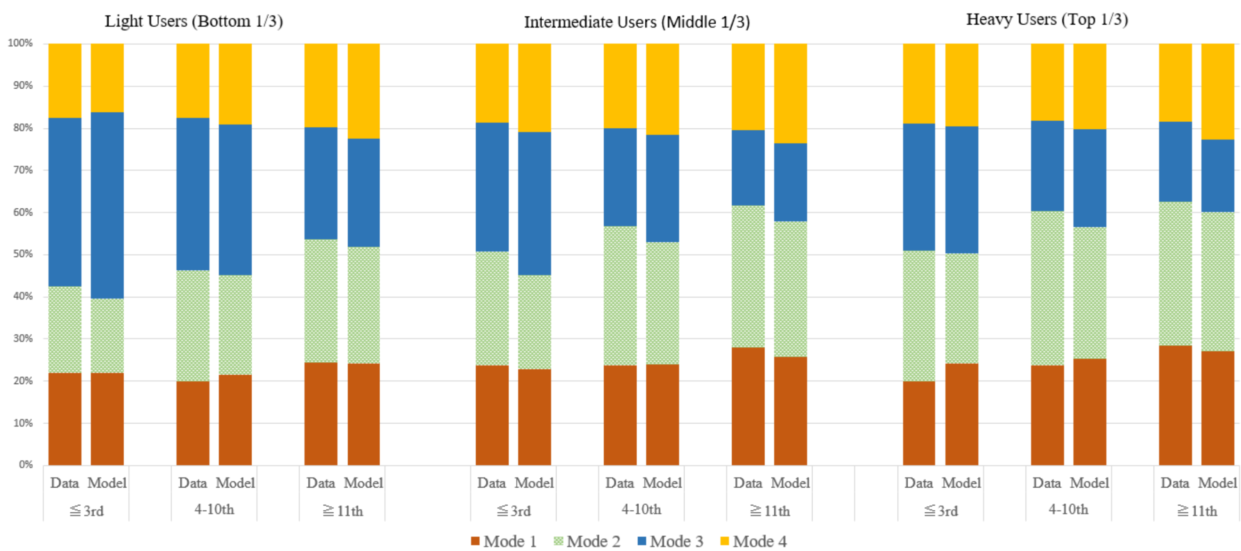

Figure 9: Model fit of game mode choice for each usage intensity

Note: The data part is identical to Figure 3. The model counterpart is computed from 50,000 simulation sequences. Usage intensity is defined as in Figure 3.

In Figure 9 through 12, I present the model fit for each of the main data variations. In Figure 9, I show the model fit for the aggregate pattern of game mode selection across users with different intensity, and its evolution over time. Heterogeneities in both cross-sectional and intertemporal dimensions are well captured. Light users tend to play mode 3 while heavy users prefer mode 2. Moreover, users gradually switch from mode 3 to other game modes. The model slightly overestimates the probability that mode 3 is selected at the beginning, and that mode 4 is selected in the long-run, but the other parts fit the data quite well.

In Figure 10, I present model prediction hit rate of each user's game mode selection. The hit rate is calculated as the choice probability that the model assigns to the mode actually selected by each user at each session, conditional on her usage history up until that point. It is obtained by integrating Equation 1 over the unobservable beliefs: $\mathbb{E}(P_m(\Omega_{it})|\{\nu_{imt}\}_{m=1}^{M}, h_t)$. In order to integrate over the distribution of the belief conditional on past actions, I use simulation with

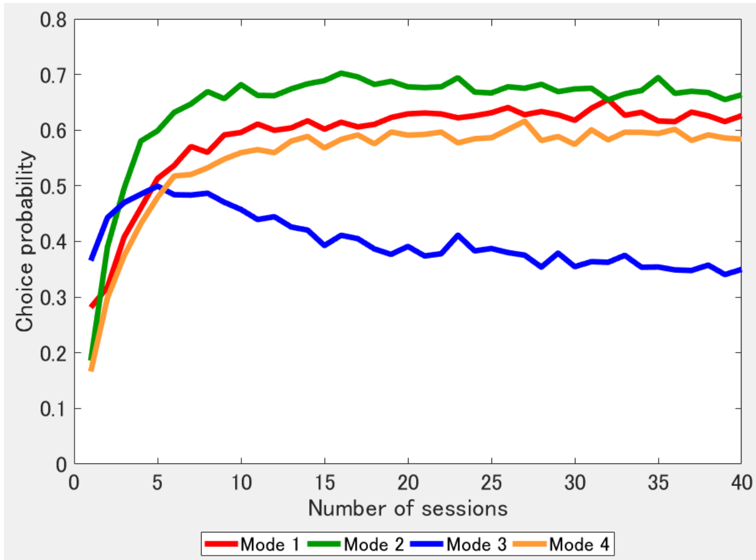

Figure 10: Model prediction hit rate: individual-level game mode selection

Note: The hit rate is the probability that the model assigns to each of the observed mode selections of each user at each session: $\mathbb{E}(P_m(\Omega_{it})|\{\nu_{imt}\}_{m=1}^{M}, h_t)$. The figure shows its average across users who selected each mode at each session.

importance sampling (Fernandez-Villaverde and Rubio-Ramirez 2007). Details of this procedure are provided in the Appendix.

Each of the lines in Figure 10 represents the model hit rate for each user at each session, averaged across users selecting the same mode. Since no history is available to be conditioned on at the beginning, the choice probability the model assigns to each user's action is close to the empirical proportion that each mode is selected. Over the first few sessions, the information of past actions significantly improves the hit rate. As experimentation ceases around the 10th session, the prediction hit rate reaches its peak at around 60 to 65 percent. On the other hand, the hit rate for model 3 remains relatively low. As shown in Figure 6, the play records of mode 3 involve more switches than the other modes in the long run; predicting future behavior is more difficult when the user switches her choice more frequently. Nonetheless, the hit rate for mode 3 is higher than the unconditional, aggregate proportion that mode 3 is selected, which is 0.243, indicating that the model still has a certain predictive power for mode 3.

The hit rate for the game mode selection, averaged across all the sessions and the modes, is 0.542. The associated positive likelihood ratio (LR+) is 3.557. Thus, the probability that the model assigns to the mode selected in the data is 3.56 times higher than the probability the model assigns to the modes not selected. Even for the worst-fitting mode 3, the average hit rate is 0.397

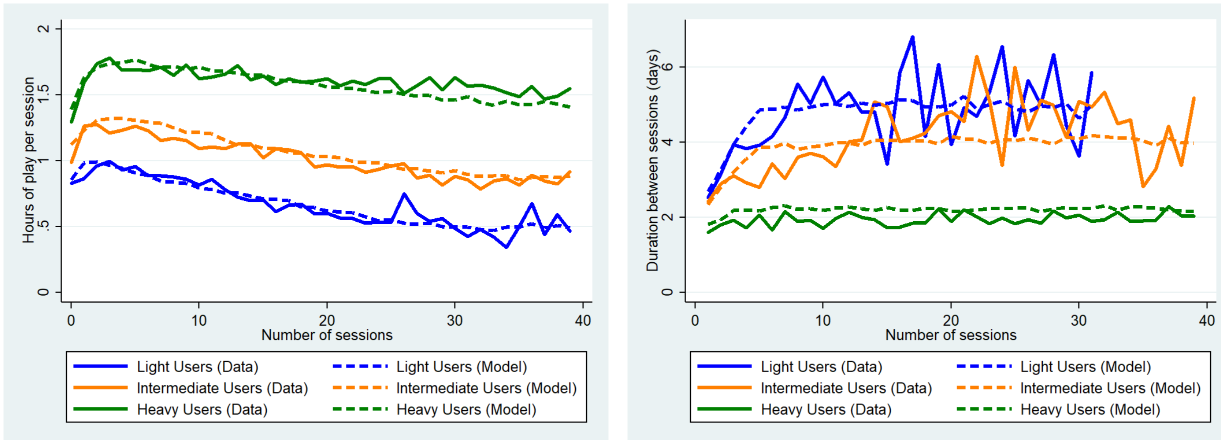

(a) The session durations

(b) The intervals between sessions

Figure 11: Model fit of the session durations and the intervals between sessions

Note: The data part is identical to Figure 4. The model counterpart is calculated from 50,000 simulation sequences. The definition of the bin is the same as in Figure 3.

and LR+ is 2.56. This is supportive evidence that the model is able to capture not only aggregate patterns, but also individual user's actions.

In Figure 11, I show the model fit of the duration of sessions and the intervals between sessions. The durations are almost perfectly captured. Since the deterministic utility evolution due to $c(\nu_{imt})$ follows monotonic decline, the initial increase of the session durations is attributed solely to the utility increase due to the reduction of the uncertainty. The interval length is also captured reasonably. Since I opt for a simple functional form for λ, the bumpy pattern of the light and intermediate users are ignored and only the average is matched. The bumpy patterns, while pronounced in Figure 11, is not correlated with users' other actions in the data. Hence, they are likely due to factors outside of the model, such as the idiosyncratic fluctuation of utility from outside options.

In Figure 12, I show the probability of termination and switching patterns. The probability of termination is underestimated around the 10th session, but the magnitude of the error is very small. The switching patterns are tracked reasonably. Two different patterns of evolution discussed in Figure 6 are both correctly matched. The estimates of other parameters and additional model fit examinations are provided in the Appendix. In particular, there I report that the model provides a reasonable fit to (1) holdout sample of 800 users, and (2) a set of users of a version released in another year.

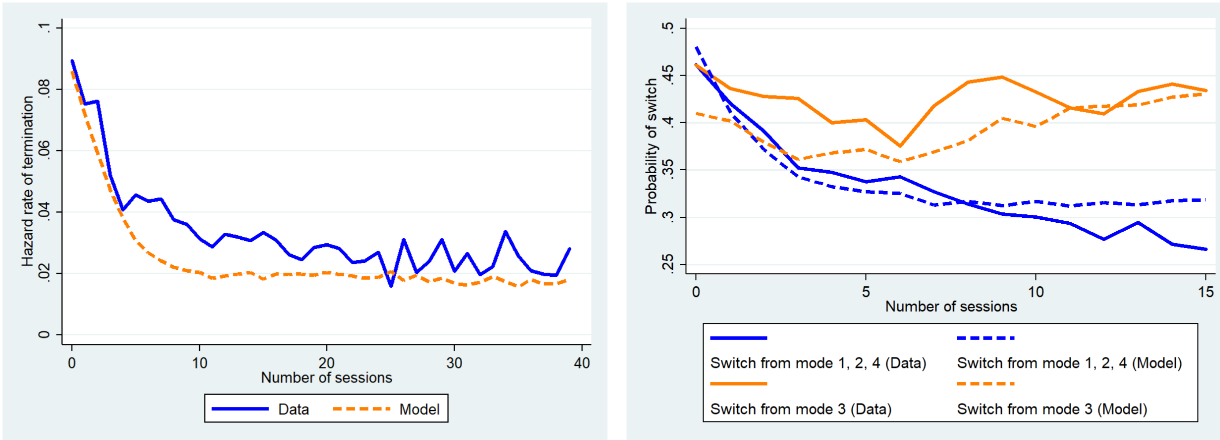

(a) The probability of termination

(b) The probability of switch

Figure 12: Model fit of termination and switching patterns

Note: The data part is identical to Figure 5 and 6. The model counterpart is calculated from 50,000 simulation sequences.

Table 3: Parameter estimates (Adoption model)

Parameters	Estimates	Std.error
Price coef mean μ_η	28.064	0.410
Price coef s.e σ_η	46.271	0.002
Arrival rate λ_u^a	0.098	0.003
Logit shock s.e σ_p	1.705	0.151

Note: Standard error is calculated by 1,000 bootstrap simulations.

6.3 Parameter estimates, model fit and implications for adoption model

In Table 3, I show parameter estimates of the adoption model. Since the price coefficient follows a lognormal distribution, the mean of the price coefficient is $\exp(\mu_\eta)$, which is at the order of 10^{12}. This is reasonable given that a vast majority of potential customers in the market don't make a purchase. In Figure 13, I show the model fit for the weekly rate of adoption up until two weeks before Christmas. The rate of adoption is defined as the number of customers making a purchase at each week divided by the total market size N. The fit is almost perfect. The existence of the initial peak and the second peak corresponding to the lower price is captured by the heterogeneity of the price coefficient.

Using the estimated parameter values, I simulate the population distribution of willingness to pay. As I only observe users who made a purchase, the population distribution is obtained by estimating a truncated distribution of willingness to pay through the method of simulated moments, and extrapolating it using the normality assumption. Figure 14 presents this entire

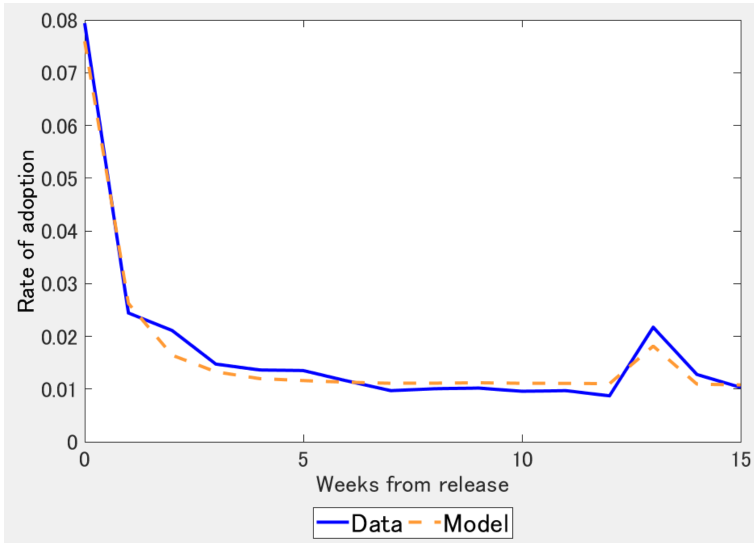

Figure 13: Pattern of purchase

Note: The data part is identical to Figure 2 except that the scale is now weekly market share. The model counterpart is calculated from 2,000,000 simulation sequences.

population distribution. The figure show the range of willingness to pay between \$1 and \$500, which covers 31.1 percent of the population. A majority of customers excluded from the figure have willingness to pay lower than \$1 and can be considered as never-buyers. The large number of low willingness to pay customers appears reasonable, for the market share of this product is 28.1 percent and majority of other customers have no intention to adopt this product. On the other hand, there exists a handful of very high willingness to pay customers: a pattern consistent with our industry knowledge.

6.4 Examining the mechanism behind customer learning

In this section, I illustrate the mechanism of learning identified in the current setting: magnitude of initial uncertainty, speed of learning and learning spill-overs, which in turn provides implications for the optimal trial design. In Figure 15a, I show the evolution of the magnitude of uncertainty that a user faces. The uncertainty is measured by the coefficient of variation of the belief: σ_{imt}/μ_{imt}. At the beginning, users face significant uncertainty. In particular, the belief of light users has a standard error that is 3.3 times higher than the mean. Reporting this in terms of willingness to pay, if a customer has an initial willingness to pay of p dollars, then the 95 percent confidence interval of her true willingness to pay is $[0.436p, 1.758p]$.[28] For example, consider a customer with

[28] Since the value function is nonlinear, the confidence interval of the willingness to pay is asymmetric despite the belief following a normal distribution.

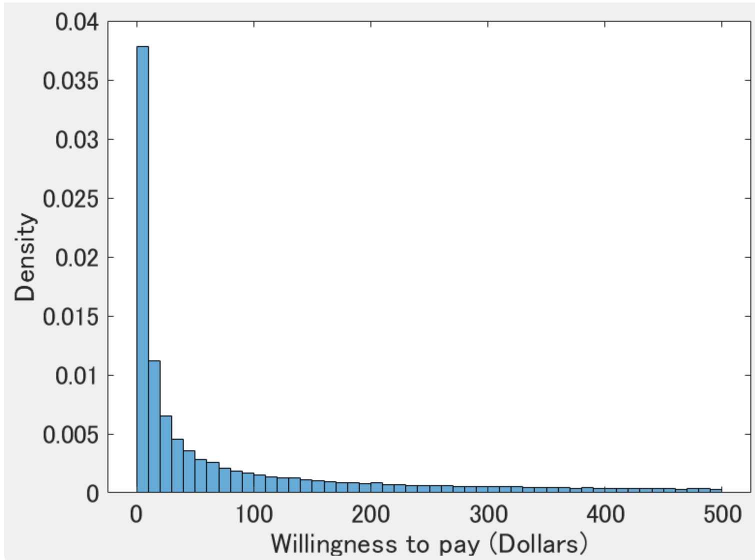

Figure 14: Distribution of willingness to pay

Note: The figure shows the distribution of willingness to pay within the interval [\$1, \$500], drawn with 2,000,000 simulated samples. Willingness to pay is defined by $V(\Omega_{i1})/\eta_i$.

willingness to pay of \$50, then her 95 percent confidence interval is [\$21.80, \$87.90]. On the other hand, heavy users face smaller uncertainty relative to their average valuation. This may indicate that users with high perceived match value engage more in pre-purchase information search. The speed of uncertainty reduction is quite fast, although the uncertainty does not collapse to zero in the short run.

The large initial uncertainty and its fast resolution indicate that willingness to pay can increase as customers accumulate usage experience. In Figure 15b, I show that indeed willingness to pay increases during the early stages of consumption. Willingness to pay reaches its peak around 5th and 6th session. The average willingness to pay after the 5th session is 13.9 percent higher than at the beginning.[29] Afterward, decrease in value due to the forgone session outweighs the informational gain and willingness to pay decreases as users play more sessions. Users with higher intensity reach the peak earlier because their valuation diminution due to the forgone session is larger than light users. Note that at the individual user level, learning does not necessarily increase product valuation. The error involved in the initial belief makes some customers overly optimistic about their match value. Those customers may be disappointed. Figure 8 shows that *on average* product valuation goes up due to the uncertainty resolution. This average increase in willingness to pay indicates that a free trial may increase demand and hence firm revenue. In particular, when the

[29]This is the average across users with initial willingness to pay between \$30 and \$60, the relevant range of users featured in the counterfactual.

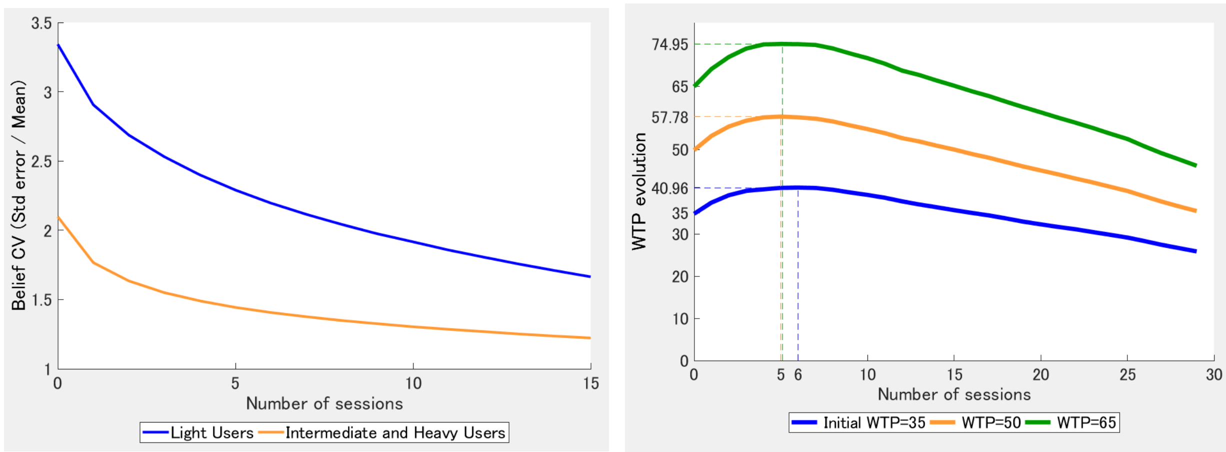

(a) Evolution of belief Std error / Mean

(b) Evolution of willingness to pay

Figure 15: Impact of learning-by-using

Note: Both panels are calculated using 50,000 simulation sequences. In Panel (a), I show the evolution of the coefficient of variation of the belief: σ_{imt}/μ_{imt}, averaged across modes and users with different usage intensity. Intermediate and heavy users exhibit almost identical patterns and I aggregate them for exposition. In Panel (b), I present the evolution of average willingness to pay for users starting from different initial values.

firm wants to provide a time-locked trial, the revenue is likely maximized around five free sessions.

In order to illustrate the magnitude of learning spill-overs, in Figure 16, I decompose the impact of learning into the effect on the belief of the selected game mode (own-effect) and that of the modes not selected (spill-over). Each of the lines represents a marginal decline in the variance of the belief due to an incremental signal received at each session. It is evident that most learning comes from the strong own-effect. One additional session decreases the variance of the own-belief by up to 63 percent, exhibiting rapid learning. On the other hand, the spill-overs play little role. The correlation of match values is not large enough for the informative signal to propagate.[30] This indicates that provision of feature-limited trials, whose profitability relies on the magnitude of learning spill-overs, may not contribute to the revenue increase in the current setup.

7 Managerial implications: the optimal trial design

Using the estimated demand parameters, I now predict how the demand responds to various trial designs and provide revenue implications. Intuitively, free trials work as follows. If a time-locked trial is provided, customers can first play any game modes and update beliefs about their match

[30]Under a Bayesian learning model with normal distribution, spill-overs are virtually non-existent for correlations below 0.8 because of the nonlinearity of the spill-over process.

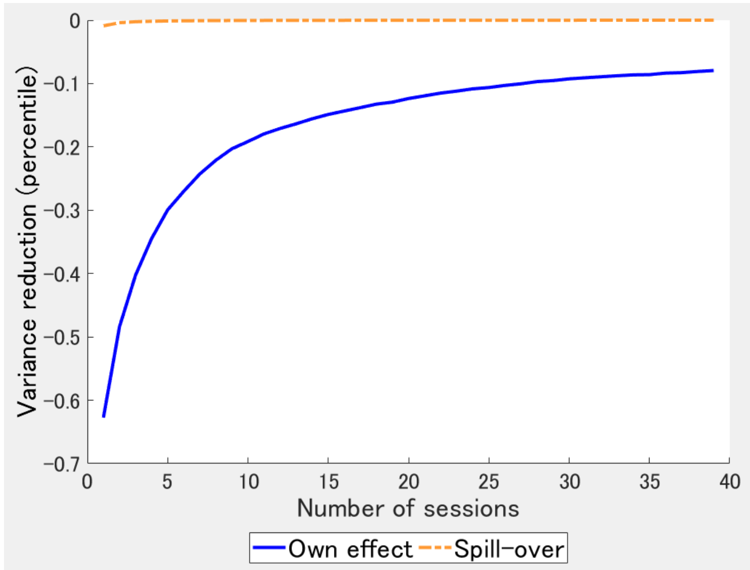

Figure 16: Marginal variance reduction from each session

Note: The figure is computed from 50,000 simulation sequences. For each i and t, I calculate the percentile reduction of the variance of the belief from t to t+1, for the mode selected at t: $\frac{\sigma^2_{im,t+1}-\sigma^2_{imt}}{\sigma^2_{imt}}$. The reported solid line is its average across the simulated sequences. The dashed line corresponds to the average of variance reduction for the modes not selected at t, calculated in a similar way.

value. Once the trial expires, customers make a purchase decision based on their posterior belief. Willingness to pay changes depending on how many sessions customers have played until they make the purchase decision. Hence, the firm's problem of choosing trial durations is equivalent to choosing at what timing customers visit the purchase occasion during the sequence of learning. On the other hand, in the case of feature-limited trials, the firm offers two vertically differentiated products: the full product and a feature-limited version. At each trial session, users can only choose features included in the trial. After the session, trial users make a purchase decision by comparing the value from switching to the full product and that from staying with the trial. They know that if they stay with the trial, they can visit another purchase occasion after the next trial session. In either trial case, in making trial usage decisions forward-looking customers take into account that learning from the current session enables them to make a better informed purchase decision in the future.

In what follows, I first describe demand models under time-locked and feature-limited trials. I then simulate customer behaviors under the estimated parameters and compare revenues under different trial designs. I refer as "time-locked trial" to the one where customers have access to the full product up to a certain number of sessions. Since I assume that learning occurs session by session, the relevant notion of a time limit is with respect to the number of sessions. Similarly, "feature-limited trial" is the one where the firm offers a subset of game modes from the full product for free. I do not consider such strategies as imposing restrictions *within* a mode; I assume that the firm only picks modes to include, and modes included in the trial remain identical to the ones in the full product.

7.1 Structural demand models under a time-locked trial

If the firm provides a time-locked trial, the firm chooses the number of free sessions, which I denote by $\tilde{T}$. Since trial includes all the features, model of trial users is similar to that of full product users up to $\tilde{T}$. After $\tilde{T}$, the trial expires and I assume that the purchase decisions thereafter are specified in the same way as in the adoption model under no trial, but with an updated belief. In addition, customers may opt to purchase the full product before $\tilde{T}$. In order to allow this, I assume that the customers also visit purchase occasions right after their arrival at the market, and also at the end of each trial session. Henceforth, I denote the optimal frequency and termination policies at each state during trial by $\tilde{\lambda}(\Omega_{i,t+1})$ and $\tilde{\delta}(\Omega_{i,t+1})$, which may be different from the post-purchase counterpart $\lambda(\Omega_{i,t+1})$ and $\delta(\Omega_{i,t+1})$. Formally, at each trial session $t \leq \tilde{T}$, a user's optimal game mode selection is specified by the following dynamic programming problem.

$$V_{it}^{TL}(\Omega_{it}, p_\tau, k_t) = \mathbb{E}[\max_{m \leq M} v(b_{it}, \nu_{imt}, h_t) + EV_{it}(\Omega_{it}, p_\tau, k_t) + \epsilon_{imt}\sigma_\epsilon], \ t \leq \tilde{T},$$

$$\text{where } EV_{it}(\Omega_{it}, p_\tau, k_t) = \begin{cases} \mathbb{E}[\tilde{\delta}(\Omega_{i,t+1})V_{i,t+1,p}^{TL}(\Omega_{i,t+1}, p_\tau, k_t) \mid \Omega_{it}, m_{it}], \text{ if } t < \tilde{T}. \\ \mathbb{E}[\tilde{\delta}(\Omega_{i,\tilde{T}+1})V_{ip}(\Omega_{i,\tilde{T}+1}, p_\tau) \mid \Omega_{i\tilde{T}}, m_{i\tilde{T}}], \text{ if } t = \tilde{T}. \end{cases}$$

The flow utility remains the same as in the full product, for all the features are included in the trial. On the other hand, the customer faces different continuation payoffs depending on whether or not they have reached $\tilde{T}$. At $t < \tilde{T}$, after the current session customers visit a purchase occasion, with their trial still remaining active. I denote the value function at the purchase occasion by $V_{i,t+1,p}^{TL}$. At $t = \tilde{T}$, the current session is the last session playable on the trial and the user makes only purchase decisions thereafter. The continuation payoff is hence identical to the value function from the model of purchase decisions V_{ip} defined earlier, but with product valuation evaluated at state $\Omega_{i,\tilde{T}+1}$. In both cases the continuation payoff is in the expectation over the realization of the signal the user receives from the current session. The state space involves two new elements. p_τ denotes the price in the current week, and $k_t \in \{1, 2, ..., 7\}$ denotes the current day within the week. These extra state variables influence the decision of the optimal timing of purchase and hence affect usage decisions through the continuation payoff. Value function V_{it}^{TL} summarizes her expected lifetime utility from adopting a free trial, including the option value of future product switches.

Now consider the purchase occasion that a user visits after session $t < \tilde{T}$. She chooses whether or not to buy the full product by comparing the value of buying to that of staying with the trial. Her value function at the purchase occasion is

$$V_{i,t+1,p}^{TL}(\Omega_{i,t+1}, p_\tau, k_t) = \mathbb{E}[\max\{V(\Omega_{i,t+1}) - \eta_i p_\tau + \epsilon_{1i\tau}\sigma_p, \bar{V}_{i,t+1}^{TL}(\Omega_{i,t+1}, p_\tau, k_t) + \epsilon_{0i\tau}\sigma_p\}], t < \tilde{T}$$

$$\text{where } \bar{V}_{i,t+1}^{TL}(\Omega_{i,t+1}, p_\tau, k_t) = \tilde{\lambda}(\Omega_{i,t+1}) \sum_{k \geq 0} (\beta(1 - \tilde{\lambda}(\Omega_{i,t+1})))^k V_{i,t+1}^{TL}(\Omega_{i,t+1}, p_{\tau+\tilde{k}}, k_t + k - 7\tilde{k}), \ \tilde{k} = \left\lfloor \frac{k_t + k}{7} \right\rfloor.$$

The value from buying includes $V(\Omega_{i,t+1})$: her valuation of the full product evaluated with the posterior belief. On the other hand, the value from staying with the trial, $\bar{V}_{i,t+1}^{TL}$, is characterized by taking an expectation of the trial value $V_{i,t+1}^{TL}$ specified above with respect to possible future price level $p_{\tau+\tilde{k}}$. $\tilde{k}$ is the number of weeks between the current and the next purchase occasion. The user takes expectations over the future prices because the date of next visit to the purchase occasion is stochastic. The next visit occurs after her next trial session, and hence the expectation is over all possible interval lengths until the next session: the same logic as in the calculation of $\beta(\Omega_{i,t+1})$ in the no trial scenario. The price is assumed to take the same value within each week τ, and exhibits a discrete jump across weeks. k_t records the day of the week at which each session t is played. If the interval between two sessions is k days, k_t evolves such that $k_{t+1} = k_t + k - 7\tilde{k}$. This results in the form of $\bar{V}_{i,t+1}^{TL}$. The solution to this dynamic programming problem provides the optimal actions and the associated value functions of a trial user at each state.

Upon arrival at the market, knowing the value of buying the full product and that of adopting a trial, the customer chooses either to adopt the trial or to buy the full product.[31] Assuming that customers arrive at the market at the first day of each week, the ex-ante value function upon the arrival at the market is described as the maximum of the two.

$$V_{i1,p}^{TL}(\Omega_{i1}, p_\tau, 1, \epsilon) = \max\{V(\Omega_{i1}) - \eta_i p_\tau + \epsilon_{1i1}\sigma_p, V_{i1}^{TL}(\Omega_{i1}, p_\tau, 1) + \epsilon_{0i1}\sigma_p\}.$$

This completes the characterization of the customer decisions under a time-locked trial.

In general, time-locked trials prompt users to defer the purchase. Since the trial product is identical to the full product until it expires, customers can play an extra trial session and reduce

[31] Since the flow utility is nonnegative and I assume that trial adoption is costless, "not buying and not trying" is weakly dominated by "trying and not buying". Hence, I do not consider the former option explicitly.

the uncertainty further by delaying the purchase at no opportunity cost. Hence, unless there is an expected price increase, the best strategy is to not buy until the trial expires. Nonetheless, in reality customers may purchase the full product before the trial expiration date. In the model, this is captured by the idiosyncratic utility shock ϵ.

Trial provision may also prompt users to experiment more with the product. Under no trial, the option value of experimenting is simply that the user can make better informed *game mode selections* in the future. In the trial, the option value also comes from better informed future *purchase decisions*. Hence, users have a higher incentive to experiment. This implies that learning is endogenous with respect to firm's policy; given the trial design selected by the firm, users choose the optimal amount of learning, changing the distribution of post-trial willingness to pay.

Other aspects of demand model not discussed here, such as the customer arrival process, are assumed to remain the same as in the no trial case. The solution to the customer's dynamic programming problem provides the probability that for a given trial restriction $\tilde{T}$, a customer with a belief b_{it} and play history $\{\nu_{imt}\}_{m=1}^{M}$ makes a purchase at price p_τ. I denote this probability by $Pr(\Omega_{it}, p_\tau, \tilde{T})$. The aggregate demand at price p_τ is equal to the probability that customers who have not purchased the product as of week τ make a purchase. It is obtained by taking the sum of $Pr(\Omega_{it}, p_\tau, \tilde{T})$ over all possible usage histories that end up with a purchase at price p_τ for all users arriving at different weeks, and integrating it over the distribution of match values.

$$D^{TL}(p_\tau, \tilde{T} \mid p_{\tau'}, \tau' < \tau) = \int \sum_{\tilde{\tau}} \sum_{\Omega_{it}, t \leq \tilde{T}} \lambda^a_{\tilde{\tau}} Pr(\Omega_{it}, p_\tau, \tilde{T}) \sum_{\tilde{\Omega}_{it}} \prod_{\Omega_{it'} \in \tilde{\Omega}_{it}} \prod_{\tilde{\tau} \leq \tau' < \tau} (1 - Pr(\Omega_{it'}, p_{\tau'}, \tilde{T})) dF(\theta_i),$$

where $\tilde{\Omega}_{it}$ represents the set of all histories that reaches Ω_{it} at period t. Because of the diminishing pool of customers, the demand at p_τ is a function of the sequence of prices that precedes it. For a given sequence of prices and for each $\tilde{T}$, one can calculate the firm revenue as

$$\pi^{TL}(\{p_1, ..., p_\tau, ...\}, \tilde{T}) = \sum_\tau p_\tau D^{TL}(p_\tau, \tilde{T}).$$

In general, choosing larger $\tilde{T}$ lets customers reduce their uncertainty further. When customers are risk averse, this increases expected utility from future sessions. However, there is an opportunity cost; initial $\tilde{T}$ sessions no longer constitute willingness to pay. In addition, some users may terminate during the trial period and the pool of customers staying in the market at $\tilde{T}$ may diminish.

7.2 Structural demand models under a feature-limited trial

Now I turn to the case of feature-limited trial. In this case, the firm chooses $\tilde{M}$ game modes to include in the trial, where $\tilde{M} < M$.[32] Assuming that customers visit a purchase occasion at the end of each trial session, trial users' dynamic programming problem is defined similarly as in the case of time-locked trial. The value function associated with the optimal game mode selection is

$$V_{it}^{FL}(\Omega_{it}, p_\tau, k_t) = \mathbb{E}[\max_{m \leq \tilde{M}} v(b_{it}, \nu_{imt}, h_t) + \mathbb{E}[\tilde{\delta}(\Omega_{i,t+1}) V_{i,t+1,p}^{FL}(\Omega_{i,t+1}, p_\tau, k_t) \mid \Omega_{it}, m_{it}] + \epsilon_{imt}\sigma_{\epsilon 1}].$$

The difference from the time-locked trial case is that the game modes available for users are now $\tilde{M}$ instead of M, whereas there is no time constraint $\tilde{T}$. The limited access to $m \leq \tilde{M}$ game modes prevents users from receiving signals from mode $m' > \tilde{M}$ and impacts the way users can update their belief. On the other hand, no time limit provides users a positive value from not buying the full product at any t. The value function at each purchase occasion is identical to the time-locked case, except for notational differences.

$$V_{i,t+1,p}^{FL}(\Omega_{i,t+1}, p_\tau, k_t) = \mathbb{E}[\max\{V(\Omega_{i,t+1}) - \eta_i p_\tau + \epsilon_{1i\tau}\sigma_p, \bar{V}_{i,t+1}^{FL}(\Omega_{i,t+1}, p_{\tau+\tilde{k}}, k_t) + \epsilon_{0i\tau}\sigma_p\}],$$

$$\text{where } \bar{V}_{i,t+1}^{FL}(\Omega_{i,t+1}, p_{\tau+\tilde{k}}, k_t) = \tilde{\lambda}(\Omega_{i,t+1}) \sum_{k \geq 0} (\beta(1 - \tilde{\lambda}(\Omega_{i,t+1})))^k V_{i,t+1}^{FL}(\Omega_{i,t+1}, p_{\tau+\tilde{k}}, k_t + k - 7\tilde{k}), \;\; \tilde{k} = \left\lfloor \frac{k_t + k}{7} \right\rfloor.$$

Upon arrival at the market, customers choose either to adopt the trial or to buy the full product.

$$V_{i1,p}^{FL}(\Omega_{i1}, p_\tau, 1) = \max\{V(\Omega_{i1}) - \eta_i p_\tau + \epsilon_{1i1}\sigma_p, V_{i1}^{FL}(\Omega_{i1}, p_\tau, 1) + \epsilon_{0i1}\sigma_p\}.$$

Intuitively, the trade-off that customers face is that by buying now, customers have full access to the product starting from the next session. On the other hand, by delaying the purchase by one more session, they can receive one more signal and reduce uncertainty at the cost of having to choose from limited features in the next session. This implies that unlike the case of time-locked trial, customers with sufficiently high initial belief prefer to buy the full product from the beginning; the negative impact from not having full access on the utility outweighs the option value from the trial. On the other hand, customers who only want the game modes provided in the trial do not benefit from buying the full product, and hence are likely to remain with the trial. The solution to

[32] In practice, the firm not only chooses the number of game modes but also *which* game mode is included in the trial. Since subscript m is just a label and has no cardinal meaning, one can always re-order modes so that the ones offered in the trial are labeled first.

the dynamic programming problem provides the probability that for a given trial restriction $\tilde{M}$, a customer with a belief and play history Ω_{it} makes a purchase at price p_τ. Denoting this probability by $Pr(\Omega_{it}, p_\tau, \tilde{M})$, I characterize the aggregate demand the firm faces at each week in the same way as in the case of time-locked trial.

$$D^{FL}(p_\tau, \tilde{M} \mid p_{\tau'}, \tau' < \tau) = \int \sum_{\tilde{\tau}} \sum_{\Omega_{it}} \lambda^a_{\tilde{\tau}} Pr(\Omega_{it}, p_\tau, \tilde{M}) \sum_{\tilde{\Omega}_{it}} \prod_{\Omega_{it'} \in \tilde{\Omega}_{it}} \prod_{\tilde{\tau} \leq \tau' < \tau} (1 - Pr(\Omega_{it'}, p_{\tau'}, \tilde{M})) dF(\theta_i).$$

The firm revenue is determined similarly.

$$\pi^{FL}(\{p_1, ..., p_\tau, ...\}, \tilde{M}) = \sum_\tau p_\tau D^{FL}(p_\tau, \tilde{M}).$$

In general, the firm wants to include features in the trial that create large learning spill-overs, so that uncertainty about features excluded from the trial also diminishes due to trial experience. At the same time the firm wants to exclude from the trial the features that many customers find high match value; the trial and the full product need to be sufficiently differentiated, in order to induce customers to make a purchase of the full product.

7.3 Simulation results

Since the aggregate demand $D^{TL}(p_\tau, \tilde{T})$ and $D^{FL}(p_\tau, \tilde{M})$ have no analytical form, I compute firm revenues at each $\tilde{T}$ and $\tilde{M}$ using 50,000 sequences of simulated customer actions. In order to highlight the main trade-offs of each trial design discussed earlier, I assume that the price is held constant at $p = 52.1$, the launch price.[33] This eliminates customers' incentive to wait for future price drops and hence any difference in the purchase timing between the case of no trial and trials is due to the incentive to learn. I also assume that the customers' optimal frequency decisions under trial remain the same as in the full product; $\tilde{\lambda}(\Omega_{it}) = \lambda(\Omega_{it})$. This is reasonable given that trial provision influences customers' usage decisions only through increasing option value from learning, and the choice of play frequency is independent of learning incentives.[34] On the other hand, as trial provision increases the option value from remaining active, users' termination rate during trial $\tilde{\delta}(\Omega_{it})$ can be lower than the no trial counterpart $\delta(\Omega_{it})$. As I do not observe free trial, I cannot estimate $\tilde{\delta}(\Omega_{it})$ from the data. Hence I simulate revenue outcome at different values of $\tilde{\delta}$.

[33] Using a different price level hardly influences the results.

[34] In this model, progression of learning is determined solely by the number of sessions; conditional on someone having played t sessions, how many days it took her to reach that state does not affect her beliefs at t.

In Figure 17, I present revenue implications of time-locked trials. The horizontal axis is $\tilde{T}$, the number of free sessions the firm provides, and the vertical axis is the percentile revenue difference from the no trial case. Each line corresponds to different value of $\tilde{\delta}$. $\tilde{\delta}$ corresponds to the speed of demand depreciation and hence the opportunity cost of providing a free trial. Two things are worth noting. First, under any $\tilde{\delta}$ such that a free trial can increase revenue, the revenue is maximized by providing five free sessions and hence it is the best trial design. This is consistent with the patterns of evolution of willingness to pay in Figure 15b. Second, revenue implications vary significantly with respect to $\tilde{\delta}$ and trials can be profitable only when $\tilde{\delta}$ is significantly lower than δ. Any trial termination rate higher than 41 percent of full product termination rate renders profitability of any time-locked trial negative. This is because significant fraction of users who terminate during the trial are indeed high willingness to pay users. For example, when $\tilde{T} = 5$, 37.6 percent of users terminating during the trial have initial willingness to pay higher than the price, and hence would likely make a purchase if there were not any trial. A trial with $\tilde{T} = 5$ is break-even when the cumulative rate of termination over five sessions is 12 percent of all users. When the rate of termination is zero, the scenario most favorable for the firm, the trial provision increases revenue by 2.54 percent. This indicates that in order to fully benefit from the trial strategy, the firm may want to incentivize users to remain active until the trial expires. In the remainder of this section, I fix $\tilde{\delta} = 0$, in order to ease the comparison across the profitability of different trial designs.

In Figure 18, I show how customers change their adoption behavior in response to the provision of a time-locked trial with five free sessions. Most of the customers whose behavior changes due to trial have original willingness to pay close to the price $p = 52.1$. This is reasonable because even a small change of the perceived match value is likely to flip the optimal action in that range. Consistent with the increase of revenue, the number of adopters increase post-trial. Considering users whose original willingness to pay is between \$30 and \$60, five trial sessions increase their willingness to pay by 15.3 percent on average. Compared to post-purchase increase of product valuation discussed in Section 6.4, the average willingness to pay after the 5th session is higher during the trial by 1.4 percentage points. Because users have a higher incentive to experiment with the product during the trial, they reach a better informed state at the end of the 5th session.

I next evaluate the revenue implication of a feature-limited trial. In the first column of Table 4, I present revenues when the firm restricts user access to only one of the game modes. The revenues are reported as a percentile difference from the no trial scenario. I find that feature-limited trials

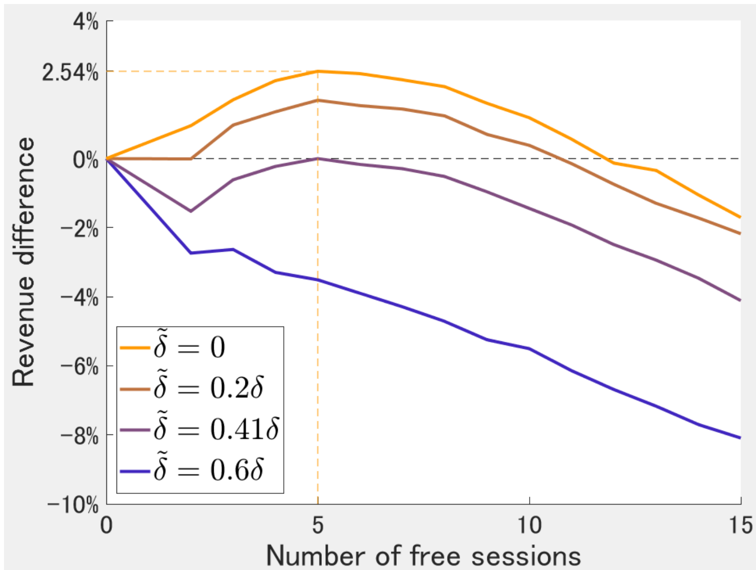

Figure 17: Revenue impact of time-locked trial provision

Note: Revenues from each trial design is calculated using 50,000 simulation sequences of users. Each line corresponds to the revenue prediction under different value of $\tilde{\delta}$. Revenues are presented as a percentage difference from no trial scenario. The price is fixed at $p = 52.1$.

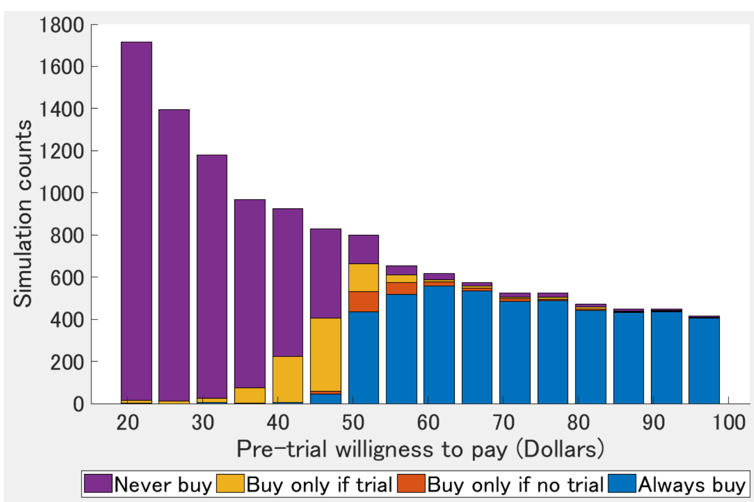

Figure 18: Customer response to trial provision

Note: The figure shows that among the customers whose original willingness to pay lies between 20 and 100 dollars, which user makes a purchase under no trial case and the case of time-locked trial with $\tilde{T}$=5. The distribution of the original willingness to pay and the users' actions are calculated from 50,000 simulation sequences. The price is fixed at $p = 52.1$. I assume $\tilde{\delta} = 0$.

Table 4: Revenue implications for feature-limited trials

	Feature-Limited only	Combined with $\tilde{T} = 5$
Mode 1 only	-12.94%	3.24%
Mode 2 only	-12.57%	2.90%
Mode 3 only	-14.82%	3.10%
Mode 4 only	-14.22%	2.99%

Note: Each cell represents the revenue from a feature-limited trial with the restriction specified by the row, measured as a percentile difference from no trial case. The revenues are calculated from 50,000 simulation sequences. The price is fixed at $p = 52.1$. I assume $\tilde{\delta} = 0$.

without any time limits do not increase revenue.[35] There are two reasons behind this. First, the product studied only contains four game modes. Hence, giving one away already sacrifices a significant portion of the product value. In general, a feature-limited trial is more profitable when a product includes more features. Second, this product does not exhibit large learning spill-overs. Providing one game mode hardly facilitates learning of match values with other modes.

Finally, I consider a situation where the firm combines both restrictions on time and feature access: a "hybrid" trial (Cheng, Li and Liu 2015). In this case, customers' dynamic programming problem is described similarly as the one from the time-locked trial, with an extra restriction that only $\tilde{M} < M$ modes are accessible during the trial. In column 2 of Table 4, I calculate revenue from adding a restriction that users can access only one mode, to the ideal time-locked trial with $\tilde{T} = 5$. The best performance is achieved when user access is limited to only mode 1. In this case, revenue increases by 3.24 percent from no trial case: 0.7 percentage point of extra revenue increase over the pure time-locked product. This is because by imposing an additional feature restriction, the firm can prompt some high willingness to pay users to make a purchase. For example, by only allowing access to mode 1, users whose most preferred mode is not mode 1 maintain high product valuation after the trial expiration and may be prompted to buy. The same set of users are less likely to buy if they are provided five free sessions of their favorite mode. The results indicate that while a feature-limited trial in itself does not make profit, it can help boost revenue from a time-locked trial by creating an additional dimension of product differentiation between the trial and the full product.

[35] While not reported, when multiple game modes are provided in the trial, the revenue is even lower.

8 Concluding remarks

In this study, I consider the impact of trial design on firm revenue. I develop a model of customer learning-by-using of a durable good with multiple features and identify the mechanism that influences trial profitability. I find that customers are risk averse and the magnitude of uncertainty around customer-product match value is large. This implies that trial provision increases customers' product valuation even when their utility declines over time due to other factors, such as boredom. I find that customers learn quickly, but learning spill-overs across different game modes are small.

This study offers several substantive insights. I find that in this setting, time-locked trials perform better and providing five free sessions is the best design, increasing revenue by up to 2.5 percent. Moreover, if the firm is willing to combine both time and feature restrictions, providing only five sessions of game mode 1 boosts revenue by an extra 0.7 percentage point. On the other hand, feature-limited trials without duration restrictions are not profitable due to small learning spill-overs. I also find that the best time-locked trial is profitable only when the cumulative rate of user termination during the trial period is less than 12 percent, which is significantly lower than the observed post-trial termination rate. This indicates that the firm may want to incentivize trial users to remain active until the trial expires. In fact, the finding is consistent with empirical observations that many videogame providers offer so-called "daily rewards" to users: a user receives a reward, such as in-game currency or items, by merely logging in to the game. The reward value increases for every consecutive day that the user logs in.

The methodology presented in this study allows firms to identify the mechanism behind customer learning from a typical data set of customers' engagement with the product. It hence helps firms in assessing the profitability of various trial designs. In particular, the structural approach does not require any observation of past trial. My model is applicable to other cases where customer learning exists and firms offer products with multiple features. Digital goods satisfy these criteria: smartphone apps and subscription services. Moreover, other goods can also exhibit similar attributes. For example, gym memberships typically entail some uncertainty (match with the instructor, offered classes, etc.) and multiple services are offered. The model helps determine whether the trial should be offered as time-locked or feature-limited (e.g. only access to yoga class).

This study contributes to the literature by offering a novel application of a Bayesian model of forward-looking customers to a durable goods setup. One nature of durable goods environment is that purchase and consumption are separable. Providing a free trial is essentially the firm's choosing when and how the customers visit purchase occasions during the sequence of learning. Moreover, learning is endogenous in this model; the firm policy directly influences how customers learn through customers' endogenous usage decisions. To my knowledge, this is the first empirical analysis that considers such interaction between customer actions and the firm policy.

This study has certain limitations when applied to other environments. I assume that the number of features in the full product is fixed and known, and there is no quality learning. Appropriate modifications are necessary in order to apply the current framework to environments where the firm introduces new features frequently, or the firm wants signal product quality through trials. Indeed, I view this study as a first step toward more general understanding of customer learning in durable goods contexts and its managerial implications. In this study, I only focus on the optimal trial design at a fixed price. The natural next step is to analyze the interaction between the optimal trial design and pricing. Although the average willingness to pay increases by 15.3 percent due to learning, under the fixed price the revenue increases by only 2.54 percent, leaving room for pricing optimization. In particular, providing a free trial enables the firm to observe each customer's usage patterns during the trial period, based on which the firm may be able to offer customized pricing. Moreover, the model I develop in this study is essentially a combination of a Bayesian model and a model of purchase timing. More generally, by offering different timing of payments the firm may be able to exploit customer learning effectively. For example, introducing subscription service or pay-per-use scheme allows customers to cancel subscription when the realized match value is low. This option value from being able to drop out can increase willingness to pay when customers are risk averse. I leave these extensions of the model for future researches.

Appendix

A Sample of customers used in this study

The sample of users consists of a set of first-time users who made a purchase of version 2014 on Sony PlayStation 3, PlayStation 4, or Microsoft Xbox360 console. "First-time users" refer to customers who had no experience with versions released in preceeding three years. This is the set of customers who had no trial access. On the other hand, some users who play this game on Microsoft Xbox One console had access to a trial version; the firm provided a 6-hour time-locked trial to customers who owned Xbox One console and subscribed to the firm's loyalty program. I opt not to use these samples, because the perfect overlap between the trial access and the enrollment to the loyalty program creates a complicated sample selection problem between the customer match value and the trial access. Rather, I use the sample of users where no such issue exists, identify customer learning and recover the effect of trial provision in a structural way.

I impose a few extra restrictions to select observations to use during estimation. First, I focus on the set of customers making a purchase within 35 weeks of product release. Since a new version of the title is released annually, customers making a purchase at later weeks may switch to the newer version and terminate the play of the older version earlier than they would without the newer version. Eliminating the last 17 weeks from the sample is sufficient, for vast majority of users terminate their play within 17 weeks. Second, in creating the moments for the adoption model I only use the data on purchases up to 16 weeks from the product release, which is two weeks before the Christmas. Customers activating the product on Christmas are likely to have received it as a gift. Hence, their activation should not be counted as a purchase when estimating the price coefficient. Since the purchase is made by the diminishing pool of customers, the number of purchase at each week is a function of the history of purchases that precedes that week. Hence, dropping Christmas period implies dropping all the post-Christmas periods as well.

B Additional evidence of customer learning in the data

Practice mode as the initial choice Aside from the four main game modes used in the analysis, the game also features an extra mode called "practice mode". In practice mode, users repeatedly conduct tasks necessary to play well in other modes. This helps new users gain an understanding about the basics of the game, and develop playing skills. On the other hand, practice mode in itself does not provide much excitement. It neither offers full matchups, nor any team-building and players' career simulation. Since users typically play practice mode only at the very beginning, I do not explicitly treat the mode as a feature and drop all practice mode sessions from the sample. However, the raw data including practice mode indicate that first-time users tend to start from pracitice mode. In Table 5, I present the proportion of game mode selection in the initial session, including practice mode. Nearly 40 percent of first-time users choose practice mode. Given the nature of practice mode, such observation implies two things. First, new users are not familiar with the game, and thus the assumption of match value uncertainty appears reasonable. Second, users do not choose to start with one of the main game modes and learn on the way. Instead, they are willing to incur the cost of forgone flow utility to gain the return, either informational or skill, in the future. In other words, there is a strong indication that users are forward-looking.

Table 5: Game mode selection and duration in the initial session

	Choice probability	Hours of play	
		Mean	Std. Dev.
Mode 1	0.134	0.603	0.819
Mode 2	0.129	0.859	1.091
Mode 3	0.245	0.3	0.595
Mode 4	0.107	0.63	0.872
Practice	0.384	0.212	0.436

Note: The table represents user behaviors in the very first session, aggregated across all users.

Initial increase of the duration of sessions In Figure 4 in Section 3, I showed that session durations increase over time in the initial few sessions. In Figure 11 I showed that the learning of risk averse customer can capture such pattern quite well; the uncertainty reduction increases the expected utility and hence users play longer. Moreover, I found that such initial increase of the hours spent is not attributable to novelty effect and skill acquisition. As I showed in Figure 8, $c(\nu_{imt})$ is monotonically increasing.

In this section I introduce three other stories that can cause the initial increase of durations, and argue that learning still plays a role even after controling for them. The first alternative story is sample selection; users with short initial duration tend to drop out earlier, and hence the duration of users who survive is longer. In Figure 4, I condition on the set of customers who remains active for 10 sessions, and hence the initial increase comes from the evolution of actions of a user and not from sample selection. The second story is the selection of game mode; users tend to choose the mode that requires less time, such as mode 3, and switch to more time-consuming modes in later

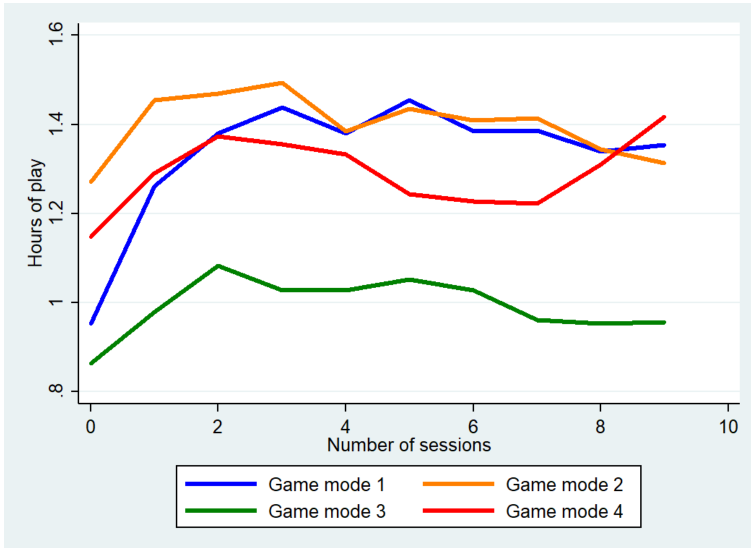

Figure 19: The average session duration of each game mode

Note: In computing this figure I use a subset of users who remain active at least until the 10th session.

periods. In order to consider such possibility, in Figure 19, I show the average session duration conditional on each game mode for users who remain active until the 10th session. The duration still increases within a mode in first few sessions, indicating that there exists a factor that increases utility within a mode.

The last alternative is the existence of day-level time constraint. If users have daily time constraint that is constant across days, and allocate the available time to more game modes at the beginning for experimenting purpose, the session duration is naturally shorter. I observe users tend to play multiple sessions per day in early periods, and hence this story applies to my data. However, as shown in Figure 20, even after aggregating the usage up to daily level, I still observe the shorter initial duration. Hence, the story of time constraint alone cannot fully explain the initial lower usage intensity.

C Model specification in detail

Structural interpretation of frequency choice and termination In the main text I assumed that the decisions of play frequency and termination are represented in a reduced form way by a probability denoted by $\lambda(\Omega_{it})$ and $\delta(\Omega_{it})$. In this section I show that a decision process where customers compare payoffs from each of the options and receive an idiosyncratic utility shock generates policy functions that are consistent with this representation.

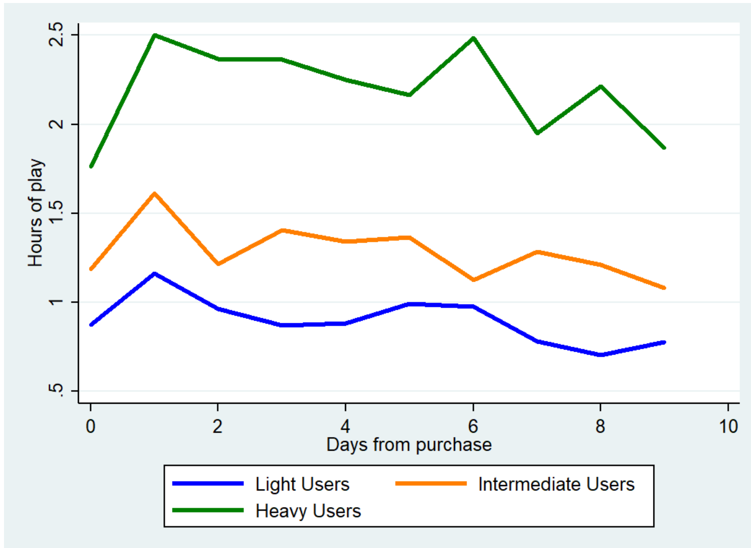

Figure 20: Average daily hours of play for each usage intensity

Note: Daily hours of play of a user is obtained as the sum of the hours spent in sessions played in each day. Presented in the figure is the average of this measure across users for each bin of usage intensity.

Consider a user's decision at the beginning of the day (node A in Figure 7). She chooses whether to play a session or not by comparing the value from playing and that of not playing. The value from playing is simply $V(\Omega_{it})$, the value function defined at node B. On the other hand, the value of not playing is computed in the following way. Suppose that if the user does not play today, then starting from tomorrow she follows a policy such that she plays a session with probability $\lambda(\Omega_{it})$ on a given day. Then the value from not playing today is the expected discounted sum of value from playing session t at some future date, where the expectation is taken over when the user plays the game next time. Denoting this expected discount factor by $\beta'(\Omega_{it})$, the value of not playing today is $\beta'(\Omega_{it})V(\Omega_{it})$, where $\beta'(\Omega_{it}) = \frac{\beta\lambda(\Omega_{it})}{1-(1-\lambda(\Omega_{it}))\beta}$.

Here I assume that the user receives an idiosyncratic utility shock for each of the options. Denoting the realization of the shock by ϵ_f, her optimal policy is defined as

$$\max\{V(\Omega_{it}) + \epsilon_{f1}, \beta'(\Omega_{it})V(\Omega_{it}) + \epsilon_{f2}\}. \tag{10}$$

If we assume that ϵ_{f1} and ϵ_{f2} follows type 1 extreme value distribution, then the user's optimal policy is represented as $\lambda(\Omega_{it}) = \frac{\exp(V(\Omega_{it}))}{\exp(V(\Omega_{it}))+\exp(\beta'(\Omega_{it})V(\Omega_{it}))}$. On the other hand, if we assume ϵ_{f1} follows Normal distribution with zero mean and variance σ^2 and $\epsilon_{f2} = 0$, then $\lambda(\Omega_{it}) = \Phi\left(\frac{V(\Omega_{it})-\beta'(\Omega_{it})V(\Omega_{it})}{\sigma}\right)$. The nonparametric representation of $\lambda(\Omega_{it})$ employed in this study encompasses these as a special case. Similar argument applies to $\delta(\Omega_{it})$.[36][37]

An extension of frequency choice In Section 5, I defined λ as the probability that a user plays the game at each day. There I assumed that λ only depends on Ω_{it}. However, this also implies that the probability of playing a session in a day does not depend on any calendar day notion, such as the number of sessions the user already played *on the same day*. In general we expect that the probability of playing another session decreases in the number of sessions played on the same day. Hence, in the empirical analysis I let the probability that "a user plays one session" and that "the user plays another session conditional on already playing at least one on the same day" be different. I denote the former as $\lambda_1(\Omega_{it})$ and the latter as $\lambda_2(\Omega_{it})$. This distinction would change the representation of the discount factor as follows.

$$\begin{aligned}\beta(\Omega_{i,t+1}) &= \delta\lambda_2 + \delta(1-\lambda_2)\lambda_1\beta + \delta(1-\lambda_2)(1-\lambda_1)\lambda_1\beta^2 + ...\\ &= \delta\left(\lambda_2 + (1-\lambda_2)\frac{\beta\lambda_1}{1-(1-\lambda_1)\beta}\right).\end{aligned}$$

Recall that $\beta(\Omega_{i,t+1})$ is located in the continuation payoff such that a user already played one session in the day. This implies that in the path of continuation, the probability that she plays the next session on the same day is always λ_2 and that she does not is $(1-\lambda_2)$. On the other hand, on the next day and after, any session she plays is always the first session of the day, and the probability that she plays a session is λ_1. During the simulation of sequences to calculate moment conditions, I use these λ_1 and λ_2 in accordance with the definition; I calculate a user's action using λ_1 if she is at the beginning of a day, and using λ_2 if she played one session on the same day. During the estimation procedure, I estimate $\lambda_1(\Omega_{it})$ and $\lambda_2(\Omega_{it})$ directly from the data.

Note on model normalization Discrete choice problems require two normalizations of utility to control for the indeterminacy of level and scale. In this model the imposed normalizations are as follows; (1) flow utility from not playing is zero, both before purchase (not purchasing) and after purchase (not playing a session), and (2) both utility from gameplay and session duration are influenced by the belief b_{it} only through $f(b_{it})$, where f is monotone with respect to μ_{imt} for a

[36] In these special cases the optimal policy only depends on Ω_{it} through $V(\Omega_{it})-\beta'(\Omega_{it})V(\Omega_{it})$. The nonparametric policy is hence consistent with more general representation of payoffs, such as inclusion of arbitrary utility from choosing outside option.

[37] Precisely speaking, this representation is inconsistent with the representation of $\beta(\Omega_{i,t+1})$ in Equation (6), where users do not take into account the realization of future ϵ_f. To the extent that those stochastic utility shocks are merely summarizing factors outside of the model, forward-lookingness with respect to such factors is not substantial.

given σ^2_{imt} and has no scaling parameters.[38] The first assumption follows the standard practice in the literature and normalizes the level of the flow utility $v(b_{it}, \nu_{imt}, h_t)$.[39] The second assumption normalizes the scale of the flow utility by that of the observed session duration. Both the session duration and the flow utility in the initial period are determined by $f(b_{it})$ and that it is monotone. Hence, there is a one-to-one mapping from the session duration to the associated utility level. Moreover, $f(b_{it})$ has no scaling parameter and hence utility has the same scale as the session duration, which is "hour". Once this assumption provides a scale of the utility and the value function, no extra normalization on the variance of the idiosyncratic shocks, both for game mode selection and for purchase, is necessary.

D Identification

The main identification challenge in this study is to separately identify the parameters of learning $\{\Sigma,\ \tilde{\Sigma}\ \sigma^2_s\}$ from other forms of state dependence $c(\nu_{it})$. The key to the identification is that the evolution of $c(\nu_{it})$ is deterministic conditional on the observed ν_{it}, while learning involves stochastic evolution of the utility. This implies that the evolution of the *variance* of the actions identify learning, while the evolution of the *mean* of the actions identify other forms of state dependence.

Formal identification of $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ at each $\bar{\Omega}_{it}$ In order to identify learning parameters $\{\Sigma,\ \tilde{\Sigma},\ \sigma^2_s\}$, I first identify $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ at each observed state $\bar{\Omega}_{it} = \{\{\nu_{imt}\}_{m=1}^{M}, h_t\}$; how the variance of the mean beliefs evolves over time. As I discussed in Section 5, observing the variance of the session duration, $Var(x^*_{imt} \mid \bar{\Omega}_{it})$, is sufficient to identify it. The argument goes as follows. In general, x^*_{imt} is a function of $b_{it} = \{\mu_{it}, \Sigma_{it}\}$ and ν_{imt} through Equation (4). However, since everyone at $\bar{\Omega}_{it}$ has the same usage history, ν_{imt} does not influence $Var(x^*_{imt} \mid \bar{\Omega}_{it})$. Moreover, Equation (9) indicates that users who share the same usage history must have the same Σ_{it} too; updating of Σ_{it} only relies on the history of choices, and not on the realization of past signals. Hence, the distribution of the session duration among users at $\bar{\Omega}_{it}$ solely reflects the distribution of their μ_{it}; there is a one-to-one mapping from $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ to $Var(x^*_{imt} \mid \bar{\Omega}_{it})$. Moreover, this mapping is monotone, and hence we can invert it to identify $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ from $Var(x^*_{imt} \mid \bar{\Omega}_{it})$. The monotonicity comes from the fact that $f(b_{it})$ is a known function for a given set of parameters and it is monotone in μ_{imt} at each $\bar{\Omega}_{it}$.

[38] In other words, if I write $f(b_{it}) = c + b * E[\theta^{\rho}_{im} \mid \theta_{im} > 0, \mu_{imt}, \sigma^2_{imt}]$, the normalization is $b = 1$.

[39] More precisely, it suffices to assume that for each of the sessions t, the utility from playing is zero at one of the state realizations.

Note that in practice, I only observe the distribution of the duration of session for the selected game modes; I observe truncated distributions and not the population-level distribution. However, the belief follows normal distribution and the point of truncation is determined by a fully parametrized model. Since normal distribution is recoverable, observation of arbitrarily truncated distribution, together with the model that specify the point of truncation, is sufficient to identify the population distribution.

Formal identification of Σ, $\tilde{\Sigma}$, and σ_s^2 from $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ The identification of $Var(\mu_{imt} \mid \bar{\Omega}_{it})$ at each observed state $\bar{\Omega}_{it}$ means that now we know how the variance of the mean belief evolves over time. This gives us sufficient information to identify the parameters that characterize the variance of the beliefs: Σ, $\tilde{\Sigma}$, and σ_s^2. In order to see this, consider $Var(\mu_{i1})$, the variance of the beliefs for all modes at the initial session, and $Var(\mu_{im2} \mid m_{i1} = m)$, the variance of the belief for mode m at $t = 2$, at the state where mode m was also selected at $t = 1$.

$$
\begin{aligned}
Var(\mu_{i1}) =& diag(Var(\mu + \Sigma(\Sigma + \tilde{\Sigma})^{-1}(\tilde{\theta}_{i0} - \mu))) \\
=& diag(\Sigma(\Sigma + \tilde{\Sigma})^{-1}\Sigma). \qquad (11) \\
Var(\mu_{im2} \mid m_{i1} = m) =& Var\left(\mu_{im1} + \frac{\sigma_{im1}^2}{\sigma_{im1}^2 + \sigma_s^2}(s_{im1} - \mu_{im1})\right) \\
=& \left(\frac{\sigma_{im1}^2}{\sigma_{im1}^2 + \sigma_s^2}\right)^2 (\sigma_m^2 + \sigma_s^2) + \left(1 - \left(\frac{\sigma_{im1}^2}{\sigma_{im1}^2 + \sigma_s^2}\right)^2\right) Var(\mu_{im1}), \qquad (12)
\end{aligned}
$$

where $diag(\cdot)$ denotes diagonal elements of the argument. σ_{im1}^2 is the $\{m, m\}$ element of Σ_1, and $Var(\mu_{im1})$ is the m-th element of $Var(\mu_{i1})$. Both Equation (11) and (12) consist of $\{\Sigma, \tilde{\Sigma}, \sigma_s^2\}$, and are not mutually collinear with respect to these parameters. Hence, these equations impose distinct restrictions on the relationship among $\{\Sigma, \tilde{\Sigma}, \sigma_s^2\}$. Similarly, the distribution of session duration for game mode m at sessions where game mode $m' \neq m$ is played at the previous session provides a restriction on the correlation between m and m'. In the same way, the variance of the belief at each state is not collinear with one another and serves as an individual constraint. Since the number of possible state realizations grows without bound as t goes up while the number of parameters is finite, one can pin down $\{\Sigma, \tilde{\Sigma}, \sigma_s^2\}$.[40]

[40] There are other variations that helps identify $\tilde{\Sigma}$. For example, $\tilde{\Sigma}$ not only determines the variance of the belief, but also determines the magnitude of the error involved in the initial belief. For example, if disproportionately large number of users buying the product early at high price play very little, it indicates that the magnitude of the error involved in the initial belief is large.

Formal identification of μ and $c(\nu_{imt})$ Identification of μ and $c(\nu_{imt})$ comes from the knowledge of average session duration $\mathbb{E}(x^*_{imt} \mid \bar{\Omega}_{it})$. Specifically, the average match value of the population μ is identified from $\mathbb{E}(x^*_{im1})$. At $t = 1$, $c(0) = 0$ and hence $x^*_{im1} = f(b_{i1})$ from Equation (4). Since f is monotone in μ_{im1}, the average duration in the initial session, $\mathbb{E}(x^*_{im1})$, identifies $\mathbb{E}(\mu_{im1})$. The identification of $\mathbb{E}(\mu_{im1})$ for each m immediately implies the identification of μ. This is because $\mu_{i1} = \mu + \Sigma(\Sigma + \tilde{\Sigma})^{-1}(\tilde{\theta}_{i0} - \mu)$ and hence $\mathbb{E}(\mu_{i1}) = \mu$. Intuitively, under rational expectation the mean of the initial belief equals that of the true match value. Over time, the average session duration evolves due to $c(\nu_{imt})$. As I discussed in Section 5, learning does not influence the evolution of the average duration. Hence, $\mathbb{E}(x^*_{imt} \mid \bar{\Omega}_{it})$ at each $\bar{\Omega}_{it}$ relative to $\bar{\Omega}_{i1}$ identifies $c(\nu_{imt})$.

Identification of other parameters Other model parameters to be identified are utility parameter α, probability of termination and play frequency $\lambda(\Omega_{it})$, $\delta(\Omega_{it})$, the distribution of price coefficient μ_η, σ^2_η, customer arrival process λ^a_τ, the variance of idiosyncratic shocks $\sigma_{\epsilon r}$, σ_p and the distribution of multiple segments ξ_r. I set the daily discount factor β at 0.999.[41]

σ_ϵ is identified by the difference between relative hours spent on each game mode and the choice probability of that mode. Consider two game modes A and B, and users on average spend 2 hours on A and 2.1 hours on B: a situation that implies that utility users receive from these modes are similar. If utility is not weighted by σ_ϵ, the choice probability has to be such that B is chosen with slightly higher probability than A. If B is selected far more often than A in the data, then it follows that σ_ϵ is low and that the size of idiosyncratic shock is very small, so that its realization hardly flips the choice even when the utility difference is modest. $\lambda_1(\Omega_{it})$, $\lambda_2(\Omega_{it})$ and $\delta(\Omega_{it})$ are identified from the distribution of termination probability and play frequency at each observed state $\bar{\Omega}_{it}$. α is identified by the difference in session durations between weekdays and weekends.

The identification of the distribution of η_i, μ_η and σ^2_η, comes from the rate of purchase at periods where the price is on a declining trend. Given that users are forward-looking, heterogeneity in the timing of adoption identifies the distribution of user patience; some users are willing to wait for price drops, while others make a purchase even when they know the future price is lower. The patience in the adoption model comes from η_i. If η_i is low, then the return from future price decline is low, so is the incentive to wait. Hence, the rate of price decline and the number of purchases

[41] Daily discount factor of 0.999 corresponds to annual discount factor of 0.95. The identification of discount factor is known to be quite difficult (Magnac and Thesmar 2002). In general it requires an exogenous factor that only influences future payoffs and not the current payoff, which the current data set does not offer.

made during that period identifies μ_η and σ_η^2.

Provided that the variations in the timing of purchase under declining price are already used to identify μ_η and σ_η^2, remaining variations to identify the customer arrival process are limited. The identification of λ_u^a relies on the number of purchases at periods where price is increasing, as well as the total market share of the product. When the price is increasing, the forward-lookingness does not play any role; the return from waiting is low. Hence, the purchase rate is solely determined by the average price coefficient of customers who remain in the market. Conditional on the distribution of price coefficient, this is a function of λ_u^a. The market share of the product is also a function of λ_u^a. Hence, I use them to identify λ_u^a.[42] σ_p serves as a residual buffer between the model prediction and the data. If $\sigma_p = 0$, then the timing of purchase is deterministic for a given willingness to pay. Hence, the proportion of customers making a purchase at each week must match with the corresponding truncated CDF of the distribution of $V(\Omega_{i1})/\eta_i$. Any difference from that identifies the magnitude of idiosyncratic shock σ_p. In practice, I did not encounter any issues in identifying parameters in the model for purchase.

Finally, even when the market consists of multiple segments of customers with different population-level parameters, the argument provided above remains valid. When multiple discrete segments exist, the distribution of match value becomes a discrete mixture of normal distributions. For a given weight ξ_r, the behavior of users corresponding to the segment assigned by ξ_r identifies the parameters for that segment. For example, suppose there exists two segments, one with low mean utility and the other with high mean utility, and the probability that a customer belongs to high segment is 0.2. Then I identify parameters associated with high and low segment from the behavior of top 20 percent and bottom 80 percent of customers, respectively. Having obtained the best fit between the data and the model prediction for a given ξ_r, ξ_r is determined to best match among them.

E Details of the estimation procedure

Step-by-step procedure of simulated method of moments Here I lay out the step-by-step procedure of conducting simulated method of moments in the current study.

[42] A flexible form of arrival process is not separately identified from the heterogeneity of price coefficient. If I pick a sequence of arrival such that "the rate of arrival at τ equals the rate of purchase at τ", then I can justify all the variation in the timing of the purchase solely by the arrival process and $\eta_i = 0$ for all i; everyone buys at the week of arrival.

1. I first pick a set of candidate parameter values Θ. In order to pick the starting value, I calculate the value of the objective function described in Section 6 at 1,000,000 random parameter values, and pick the smallest one.

2. Given Θ, I solve the model of usage and purchase described in Section 4 for each segment r. I solve for the value functions by backward induction. In order to ease the computational burden, I use the discretization and interpolation method (Keane and Wolpin 1994). For each session t, I randomly pick 15,000 points from the state and evaluate the value function at these points. I then interpolate the values at other points by fitting a polynomial of the state variables. The variables included as regressors are as follows: $\mu_{imt}, \nu_{imt}, \exp(\mu_{imt}/100), \mu_{imt} \times \nu_{imt}, \mu_{imt} \times \mu_{im't}$, h_t, $\mu_{imt} \times h_t$ for all m and $m' \neq m$. Before estimating the full model, I tested the validity of the polynomial approximation by solving both the approximated value function and the full solution of a 10-period dynamic programming problem. The approximation was quite close to the full solution, with R^2 being over 0.97 at every t.

3. Once I obtain the value function at each state for each segment r, I first create moments for the purchase model. I draw 2,000,000 individuals, each of whom belongs to segment r with probability ξ_r, which is the proportion of the segment in the market. I draw their true match value θ_i following $N(\mu_r, \Sigma)$ and the initial signal $\tilde{\theta}_{i0}$ from $N(\theta_i, \tilde{\Sigma})$, and create the initial belief μ_{i1} and Σ_{1r}. I also draw their timing of arrival at the market according to λ_u^a. Using the initial belief and the timing of arrival, I simulate their purchase decisions following the policy function computed in the model of purchase. The simulated purchase decisions are used to create moments of the pattern of purchase, which are matched with the data.

4. In order to create moments for usage model, I take the subset of simulated individuals who are predicted to make a purchase in the previous step. This is the sample of users comparable to the real data. For this sample of users, I draw a sequence of usage. For each user and for each session, given the drawn state Ω_{it} I draw her actions according to the policy functions, and draw a signal realization. Using the signal I update her beliefs and calculate the state $\Omega_{i,t+1}$. The sequences of actions obtained this way are used to create moments of the usage patterns.

5. I calculate the value of the objective function and update the parameters. I repeat these steps until the convergence is achieved.

6. In order to check if the global minimum is attained, I conduct this exercise with multiple starting values.

Construction of moments From the usage model, I use the following set of moments.

1. The average and the variance of the session duration, the choice probability of each game mode, and the probability of termination, evaluated at each history of game mode selection and cumulative hours of play up to the previous session: $\mathbb{E}(x_{imt} \mid \tilde{\nu}_{it}, \bar{x}_{it})$, $Var(x_{imt} \mid \tilde{\nu}_{it}, \bar{x}_{it})$, $\mathbb{E}(m_{it} \mid \tilde{\nu}_{it}, \bar{x}_{it})$ and $\mathbb{E}(term_{it} \mid \tilde{\nu}_{it}, \bar{x}_{it})$ where $\tilde{\nu}_{it} = \{\{\nu_{imt'}\}_{m=1}^{M}\}_{t'=1}^{t}$ and $\bar{x}_{it} = \frac{\sum_{t' \leq t-1} \sum_m x_{imt'}}{t-1}$. $term_{it}$ is termination indicator, which equals one if user i terminates after session t. (868 moments)

2. The probability that the game mode selected in the next session is different from the one at the current session, and the average interval length between the current and the next session, evaluated at each history of game mode selection and cumulative hours of play up to the current session: $Pr(m_{i,t+1} \neq m_{it} \mid \tilde{\nu}_{it}, \tilde{x}_{it})$ and $\mathbb{E}(d_{it} \mid \tilde{\nu}_{it}, \tilde{x}_{it})$, where $\tilde{x}_{it} = \frac{\sum_{t' \leq t} \sum_m x_{imt'}}{t}$. d_{it} denotes the interval length between session t and t+1. (188 moments)

3. The average session duration, the choice probability of each game mode, and the average interval length between the current and the next session, evaluated at each of the cumulative number of past sessions and the cumulative lifetime hours of play: $\mathbb{E}(x_{imt} \mid t, X_i)$, $\mathbb{E}(m_{it} \mid t, X_i)$ and $\mathbb{E}(d_{it} \mid t, X_i)$, where $X_i = \frac{\sum_{t' \leq \tilde{t}_i} \sum_m x_{imt'}}{\tilde{t}_i}$. $\tilde{t}_i$ is the number of sessions user i played until she terminates. (5,250 moments)

4. The probability of termination evaluated at each of the cumulative number of past sessions and the cumulative hours of play in the initial five sessions from purchase: $\mathbb{E}(term_{it} \mid t, X_{wi})$, where $X_{wi} = \frac{\sum_{t=1}^{5} \sum_m x_{imt}}{5}$. (900 moments)

5. The probability that the game mode selected in the next session is different from the one at the current session, evaluated at each of the cumulative number of past sessions and the game mode selected in the current session: $Pr(m_{i,t+1} \neq m_{it} \mid t, m_{it})$. (120 moments)

6. The average session duration, evaluated at each of the cumulative number of past sessions and weekend indicator: $\mathbb{E}(x_{imt} \mid t, h_t)$. (60 moments)

7. The probability that users play multiple sessions within a day, evaluated at each of the cumulative number of past sessions: $Pr(1\{d_{it} = 0\} \mid t)$. (30 moments)

In order to condition the moments on continuous variables (e.g. cumulative hours of play in the past sessions), I create 10 bins for each of them and compute conditional expectations in each bin. In addition, I create moment 3 and 4 for each subset of samples who survived at least 5 sessions, 10 sessions and 20 sessions. For some users at some sessions, the record of the session duration is missing. I exclude those sessions from the calculation of moments involving session durations. The missing duration is simply due to the technical difficulty of keeping track of the timestamp of play, and no systematic correlation between the pattern of missing data and the pattern of usage was observed. Hence, it does not introduce any bias in the estimates. I only use the first 30 sessions as the moments to ease the computational burden. As shown in Section 3, most of consumption dynamics, and hence the implied customer learning, stabilize within the first 10 sessions. Hence, the variation from the initial 30 sessions is sufficient to identify both the mechanism behind learning and the distribution of the true match value.

The moments used to identify the adoption model are the rate of adoption at each week from week 1 through 16. In the model, the rate of adoption is calculated by the number of simulation paths making a purchase at each week, divided by the number of total simulation paths. In the data, the rate of adoption corresponds to the proportion of customers making a purchase at each week in the data, multiplied by the market share of the product, whose derivation is described below.

Note that the moment conditions closely follow the identification argument provided earlier. For example, the evolution of belief is identified by the average duration in the initial session and the evolution of the variance of the session duration. This is accounted for through moment 1. The identification of the coefficient of risk aversion ρ comes from initial switching pattern, which is accounted for by moment 2. In general, the evolution of the behaviors across states are the identifying variations of the parameters, and hence all the moments are conditioned on the finest possible bins of histories that maintain a certain number of observations in each bin.

Derivation of market share Here I describe the derivation of market share of the product, which is used in computing the empirical rate of adoption. I assume that the total market size for the sports games is proportional to the share of sports games among all the videogame software sales. The average share of sports games between 2007 and 2015 is 16 percent. The total market size of all games is assumed to be equal to the installed base of PlayStation 3, PlayStation 4 and Xbox

360, which is 99.42 million units.[43] Hence, the market size for sports games is $99.42 \times 0.16 = 15.91$ million. This number corresponds to N in the current study. The sales of the focal game that is compatible to the above consoles are 4.47 million units. Therefore, the market share of this title is $4.47/15.91 = 0.281$, which I use as the market share of the game.

Parametrization of f, c and λ In this section, I provide details of parametrization employed in the current study. In the main text, I assumed that f is parametrized as follows.

$$f(b_{it}) = E[\theta_{im}^{\rho} \mid \theta_{im} > 0, b_{it}].$$

In practice, I find that when ρ is very small, f defined as such is almost flat with respect to perceived match value, creating significant computational slowdown. In order to address this, I use a transformed version of it, specified as follows.

$$f(b_{it}) = (E[\theta_{im}^{\rho} \mid \theta_{im} > 0, b_{it}]P(\theta_{im} > 0 \mid b_{it}))^{\frac{1}{\rho}}.$$

This transformation makes the computation faster by an order of magnitude. Since this functional form does not have a closed form, in practice I compute $f(b_{it})$ for each ρ, μ and σ^2 using Gauss-Legendre quadrature.

c is specified as a quadratic function of the past number of sessions as follows.

$$c(\nu_{imt}, b_{it}) = (\gamma_1 - \gamma_2 f(b_{it}))\nu_{imt} + (\gamma_3 - \gamma_4 f(b_{it}))t - \gamma_5 \nu_{imt}^2.$$

In order to capture the observed pattern that the the evolution of usage intensity is heterogeneous, I allow the coefficients to depend on the perceived match value through $f(b_{it})$. Since allowing c to be a flexible function of b_{it} introduces an identification issue, I assume that c depends on b_{it} only through $f(b_{it})$, thereby maintaining identifiability.[44]

[43] Xbox One is excluded from the market size because the customers with Xbox One are excluded from the current analysis.

[44] The identification issue arises when both f and c are nonparametric. In practice, I impose a particular parametric form on f and hence having more flexible c function as a function of b_{it} is likely to maintain identification.

Table 6: Parameter estimates of c, λ and δ

Parameters		Estimates	Std.error	Parameters		Estimates	Std.error
$c(v_{it})$	γ_1	0.042	0.003	$\lambda(\Omega_{it})$	φ_{l1}	0.456	0.001
	γ_2	0.025	0.005		φ_{l2}	0.085	0.011
	γ_3	0.036	0.004		φ_{l3}	0.065	6*10^(-5)
	γ_4	0.072	0.001		φ_{l4}	0.052	0.001
	γ_5	0.005	0.0002		φ_{l5}	0.015	0.002
					φ_{l6}	-0.175	0.011
$\delta(\Omega_{it})$	φ_{d1}	0.896	0.0001				
	φ_{d2}	0.024	0.001				
	φ_{d3}	-0.017	5*10^(-8)				
	φ_{d4}	-0.008	2*10^(-8)				

Note: Standard error is calculated by 1,000 bootstrap simulations.

Similarly, $\lambda_1(\Omega_{it})$, $\lambda_2(\Omega_{it})$ and $\delta(\Omega_{it})$ are parametrized as follows.

$$\lambda_1(\Omega_{it}) = \phi_{l1} + \phi_{l2}\frac{(\bar{\mu}_{it} - \bar{\mu})}{\bar{\sigma}} - \left(\phi_{l3} - \phi_{l4}\frac{(\bar{\mu}_{it} - \bar{\mu})}{\bar{\sigma}}\right) t + \phi_{l5}t^2,$$

$$\lambda_2(\Omega_{it}) = \lambda_1(\Omega_{it}) + \phi_{l6},$$

$$\delta(\Omega_{it}) = \phi_{d1} + \phi_{d2}\frac{(\bar{\mu}_{it} - \bar{\mu})}{\bar{\sigma}} - \phi_{d3}t + \phi_{d4}t^2,$$

$$\text{where } \bar{\mu}_{it} = \frac{\sum_m \mu_{imt}}{M},\ \bar{\mu} = \frac{\sum_m \mu_m}{M},\ \bar{\sigma} = \frac{\sum_m \sigma_m}{M}.$$

Both λ and δ are quadratic with respect to the number of past sessions t, and its intercept and slope depends on the current belief. The term $\frac{(\bar{\mu}_{it}-\bar{\mu})}{\bar{\sigma}}$ represents the normalized location of the belief of user i relative to the average belief of the population. This specification allows for a possibility that a user with higher perceived match value plays the game more frequently and has lower probability of termination, and she is increasingly so as she accumulates more experience. λ_2, the probability that a user plays multiple sessions in a day, is different from λ_1 only by an additive constant. As discussed in the previous section, this extra constant term captures that even heavy users do not often play multiple sessions within a day. In Table 6, I present the parameter estimates of the functions presented in this section.

F Other figures of model fit

In Figure 21, I show the evolution of the probability that each game mode is selected. Unlike Figure 9, this shows the choice probability at every single session, while not conditional on usage intensity.

The choice pattern is tracked quite well. The choice probability of game mode 2 and 3 (1 and 4) are slightly underestimated (overestimated), but the magnitude of the error is small. In Figure 22, I show the histogram of the duration of the very first session. The first session is chosen merely for expositional purpose and the fit for the other sessions are similar. Notably, the model tracks the shape of the distribution flexibly, even though the belief follows normal distribution. This is because of the existence of multiple types. The low segment creates a mass below 1 hour, and the high segment creates one around 2 hours.

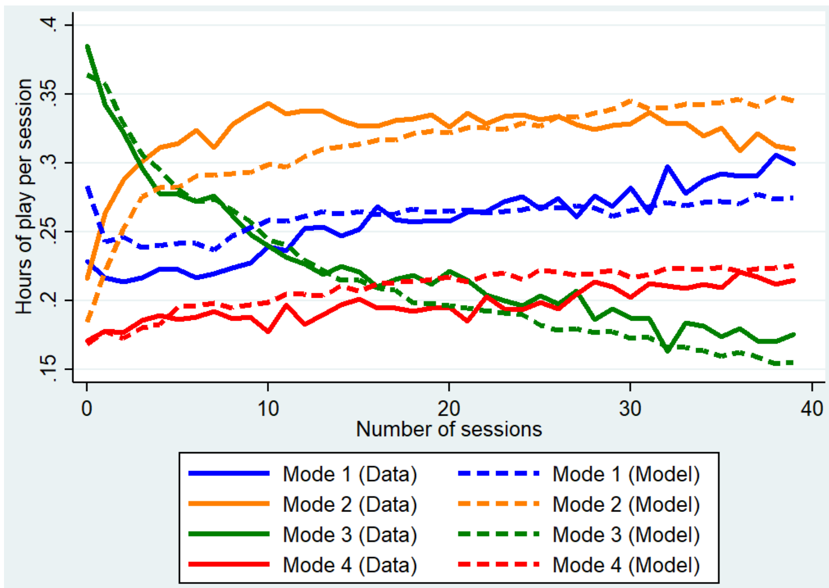

Figure 21: The model fit of the game mode selection

Note: The choice probability is computed by the number of users who select each game mode at each session, divided by the number of users who remain active. The model counterpart is calculated using 50,000 simulation sequences.

G Model validation exercises

Model fit to holdout sample Among 4,578 users in the data, a randomly selected 800 users are not used in the estimation and serve as holdout sample. In Figure 23, I present the model fit with this sample. In order to ease comparison, the figures presented here are identical to the ones presented earlier, except that the data part is replaced by that of the holdout sample. The model maintains a good graphical fit to most of the data patterns. Adoption pattern presented in Figure 23f fits less well due to the existence of secondary peak around the 5th week that does not exist in the estimation sample. On the other hand, all usage patterns exhibit a reasonable fit. The average

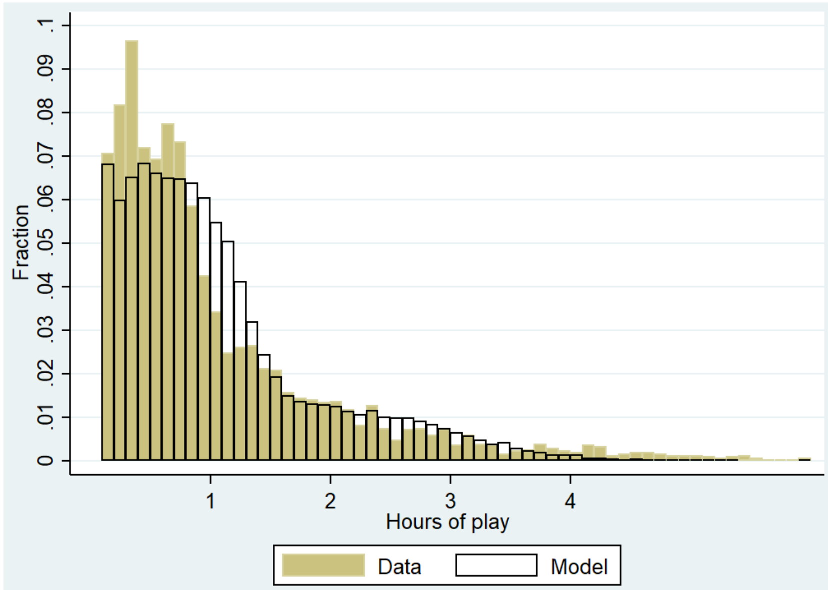

Figure 22: The model fit of the distribution of duration of the initial session

Note: Users whose duration is less than 6 minutes are dropped.

prediction hit rate for the game mode selection is 0.534, and LR+ is 3.438. Both of them are very close to corresponding ones from the estimation sample, which is 0.545 and 3.587, respectively. The model also maintains good out-of-sample predictive power for the session duration and the intervals between sessions. The ratio of standard error of prediction errors between out-of-sample and in-sample is 1.041 for the session duration, and 1.033 for the intervals between sessions.[45] In other words, the magnitude of errors involved in the out-of-sample prediction of the durations and the intervals is only 4.1 and 3.3 percent higher. They indicate that the model estimates capture the underlying mechanism common across all users, rather than merely reflect some particular variation of the estimation sample.

Model fit to users from another year Throughout the paper, I use a set of customers making a purchase of version released in 2014 as an estimation sample. Another model validation exercise is to ask whether the model estimated as such can predict actions of users from some other years. While the quality of graphics and the real-league data contained in the game are updated every year, main features mostly stay the same across versions. Hence, it is reasonable to expect similar customer behaviors across years. In order to explore this, I compare the model prediction with a set of first-time customers making a purchase of version 2015. This set of users serves as

[45] The prediction error of the duration of user i at session t is given by $\tilde{x}_{it} - \mathbb{E}(x^*(\Omega_{it}) \mid \{\nu_{imt}\}_{m=1}^{M}, h_t)$, where $\tilde{x}_{it}$ is the observed session duration. The error for intervals between sessions is defined similarly.

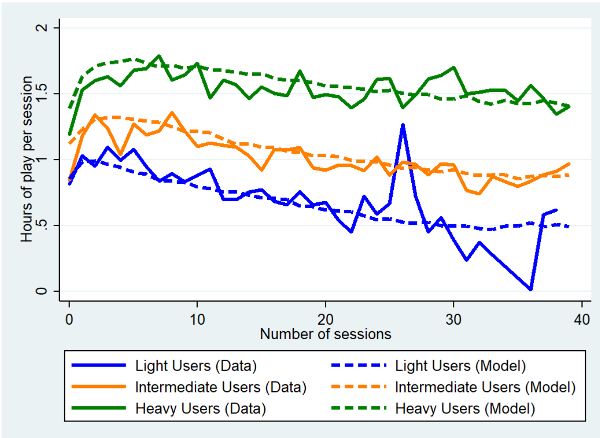

(a) Session duration

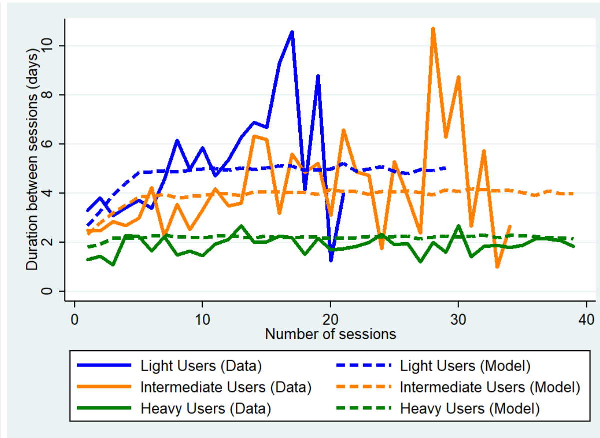

(b) Intervals between sessions

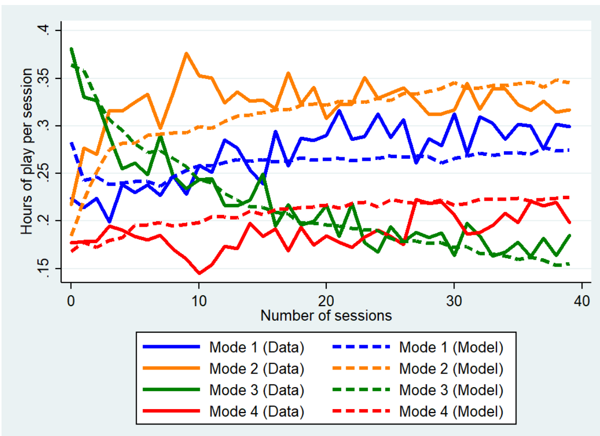

(c) Game mode selection

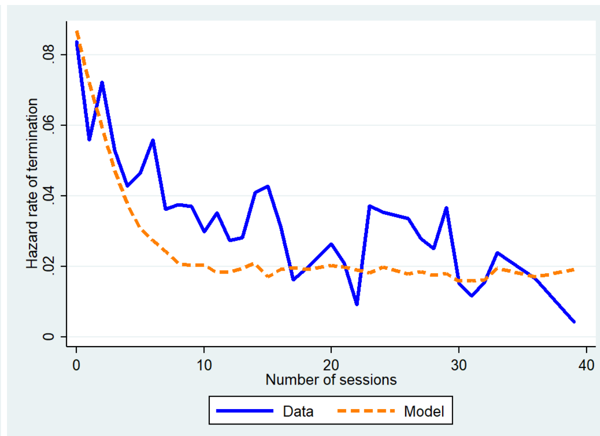

(d) Probability of termination

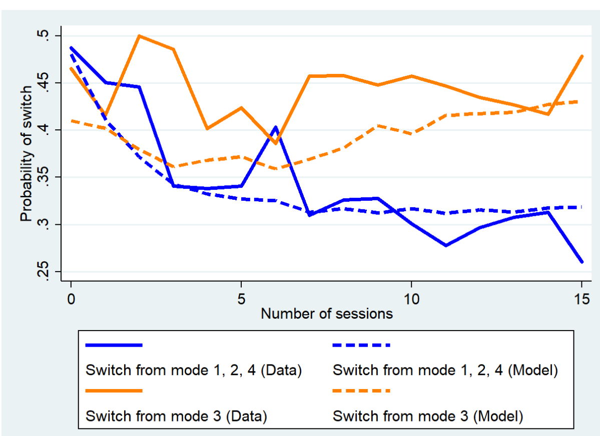

(e) Probability of switching

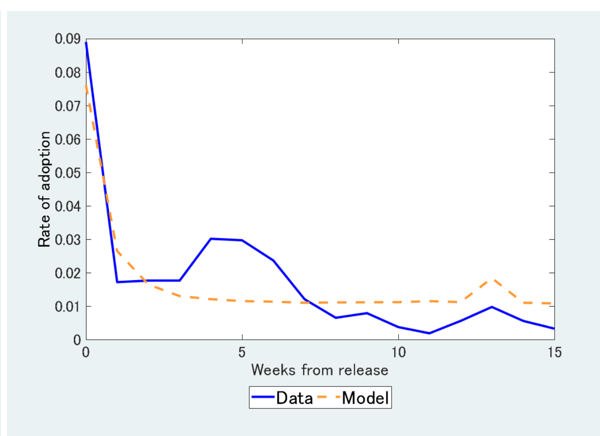

(f) Adoption pattern

Figure 23: Model fit to holdout sample

Note: The data part is calculated using holdout sample of 800 users. The model counterpart is identical to the figures presented before.

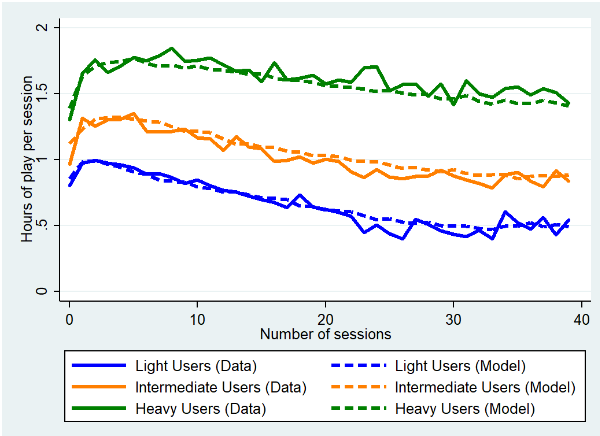

(a) Session duration

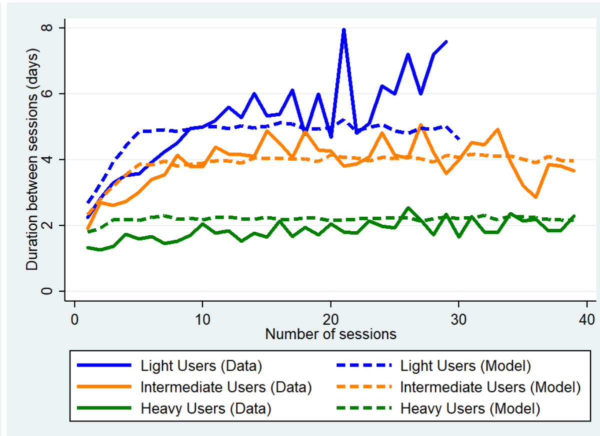

(b) Intervals between sessions

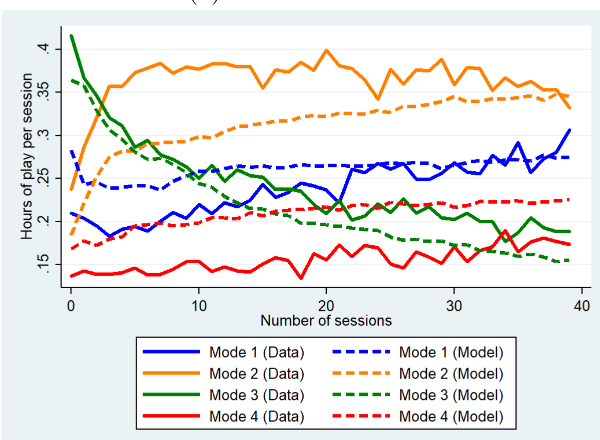

(c) Game mode selection

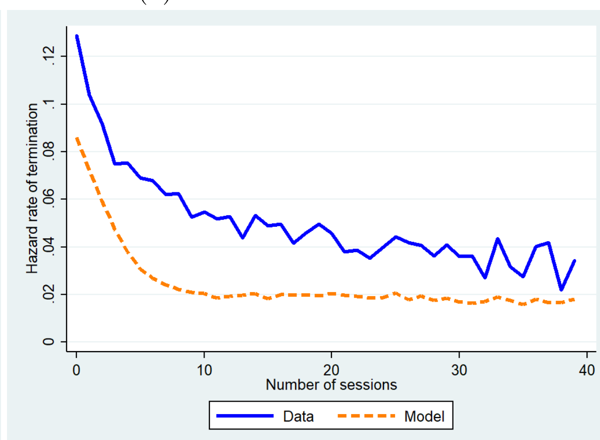

(d) Probability of termination

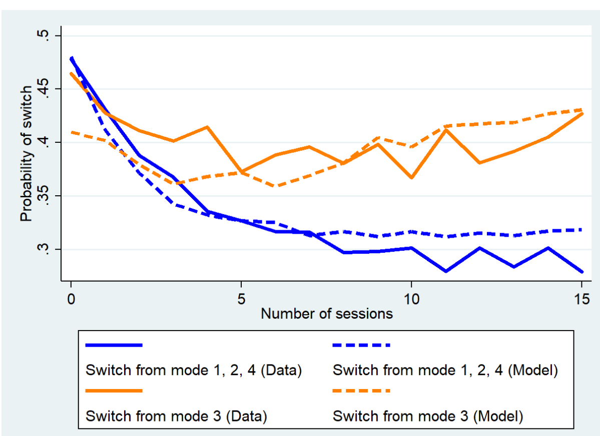

(e) Probability of switching

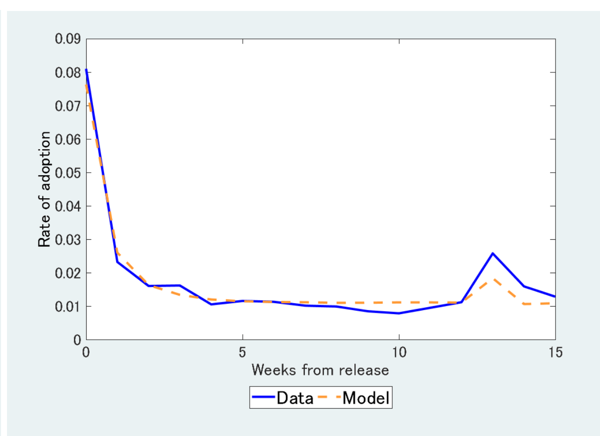

(f) Adoption pattern

Figure 24: Model fit to user actions from another year

Note: The data part is calculated using 5,211 first-time users activating version 2015. The model counterpart is identical to the ones presented in the main text.

an ideal holdout sample. Version 2015 features exactly the same number of game modes with the same name, allowing one to calculate the same measure as for version 2014. Moreover, the set of first-time users of version 2014 and that of version 2015 are mutually exclusive, providing an opportunity for pure out-of-sample fitting exercise.

In Figure 24, I present the model fit with this sample. Overall graphical fit is surprisingly good. In particular, the fit of the session durations presented in Figure 24a is almost as good as that of the estimation sample. The pattern of game mode selection presented in Figure 24c is less ideal. This is reasonable because specific characteristics of each mode provided in version 2014 and 2015 can be slightly different from each other. It is also notable that the probability of termination is in general higher in version 2015, as represented in Figure 24d. Although this may indicate the existence of possible quality issue for version 2015, such speculation is completely outside the model. Finally, the good fit of adoption pattern presented in Figure 24f is remarkable given that the model prediction is calculated using the history of prices for version 2014. This indicates both the demand structure and the price pattern are quite similar between these two years. The average prediction hit rate for the game mode selection is 0.533, and LR+ is 3.431. These are very close to the in-sample ones. The ratio of standard errors of the prediction errors between out-of-sample and in-sample is 1.025 for the session duration, and 0.977 for the intervals between sessions. In other words, the errors involved in the out-of-sample prediction is 2.5 percent higher for the session durations, and 2.3 percent *lower* for the intervals. Overall, these results are indicative that the model is not merely useful to explain user behaviors from the specific version of the product I study, but also capture a universal tendency underlying the customer behaviors.

H Importance sampling simulation to forecast individual usage pattern

In order to evaluate model prediction for each user's actions, it requires an expectation of the action specified by the model over the unobserved belief, conditional on usage history. For example, prediction of a user's game mode selection is given by $\mathbb{E}(P_m(\Omega_{it})|\{\nu_{imt}\}_{m=1}^{M}, h_t)$. I calculate this integral through simulation. In general, simulating a conditional expectation requires a sufficiently large number of simulation draws that satisfy the conditioning requirement. This is because the simulated expectation is nothing but the average of such samples. However, as we simulate a sequence of actions with a long horizon, the number of possible histories increases quickly, making it difficult to secure sufficient sample size at each state.

In order to deal with this issue, I employ importance sampling approach suggested by Fernandez-Villaverde and Rubio-Ramirez (2007).[46] The idea is that for each i and at each period t, I replace simulation sequences that do not explain the observed actions at t very well with ones that do it better. By repeating this replacement at every t, when one evaluates the conditional expectation at t+1, all sequences in the pool are likely to have the history that user i actually follows. Hence, the pool consists of more sequences that satisfy the conditioning requirement.

Formally, the simulation proceeds as follows. Suppose that I intend to calculate $\mathbb{E}(P_m(\Omega_{it})|\{\nu_{imt}\}_{m=1}^{M}, h_t)$ for all t. For each individual user in the data, I first draw her true type θ_{is} from the population distribution of match value, draw her initial signal $\tilde{\theta}_{i0s}$ and calculate the initial belief b_{i0s}. Subscript s denotes each simulation draw. At each session t, given the drawn belief b_{its} and other relevant state the model provides the probability that the mode user i selected in the data is selected. I denote this probability by $P_m(b_{its}, \{\nu_{imt}\}_{m=1}^{M}, h_t \mid \theta_{is})$. By taking its average over simulation draws, I have an estimator of $\mathbb{E}(P_m(\Omega_{it}) \mid \{\nu_{imt}\}_{m=1}^{M}, h_t)$ at t.

Moving on to period t+1, in order to compute $P_m(b_{i,t+1,s}, \{\nu_{im,t+1}\}_{m=1}^{M}, h_{t+1} \mid \theta_{is})$, it requires $b_{i,t+1,s}$: a set of draws corresponding to the belief at period t+1. While crude frequency estimator suggests that I simply draw a set of signals s_{its} for the chosen action and update b_{its} to get $b_{i,t+1,s}$, importance sampling inserts an additional step; I replace sequences that exhibits low likelihood of explaining the user's session t action with the ones with high likelihood. Specifically, I first weight each of the draws b_{its} by $\frac{P_m(b_{its},\{\nu_{imt}\}_{m=1}^{M},h_t|\theta_{is})}{\sum_{s'} P_m(b_{\theta its'},\{\nu_{imt}\}_{m=1}^{M},h_t|\theta_{is'})}$. This weight corresponds to how well each draw b_{its} explains the behavior at period t, relative to other draws $b_{\theta its'}$. The draw that fits the data well at period t receives higher weight. I then re-draw the set of beliefs from this re-weighted set of b_{its} in the same way as boot-strapping with replacement. Those with high weight may get drawn multiple times, while those with low weight may not get drawn at all. This re-draw provides a re-weighted set of b_{its}, from which I construct $b_{i,t+1,s}$ in the same way as in the crude estimator.[47] This additional step guarantees that at each t+1, the belief $b_{i,t+1,s}$ is drawn such that b_{its} explains the action taken at session t well. Hence, it is more likely that many of those sequences satisfy the conditioning requirement $\{\{\nu_{imt}\}_{m=1}^{M}, h_t\}$ for all t.

[46] In the original paper, this method is called "Particle filtering".

[47] Once b_{its} is replaced by a new sequence, the corresponding true type θ_{is} is replaced as well, so that each re-drawn b_{its} has a correct corresponding θ_{is}.

References

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

[2] Arora, Sandeep, Frenkel ter Hofstede, and Vijay Mahajan (2017), "The Implications of Offering Free Versions for the Performance of Paid Mobile Apps," Journal of Marketing 81(6), 62-78.

[3] Bhargava, Hemant K. and Rachel R. Chen (2012), "The Benefit of Information Asymmetry: When to Sell to Informed Customers?," Decision Support Systems 53(2), 345-356.

[4] Brice, Andy (2009), "The Truth about Conversion Ratios for Downloadable Software," Available from: https://successfulsoftware.net/2009/04/23/the-truth-about-conversion-ratios-for-software/

[5] Chamberlain (1984), "Panel Data," Handbook of Econometrics 2, Griliches Z, Intriligator MD, eds. 1247-1318.

[6] Che, Hai, Tülin Erdem and T. Sabri Öncü (2015), "Consumer Learning and Evolution of Consumer Brand Preferences," Quantitative Marketing and Economics 13(3), 173-202.

[7] Chellappa, Ramnath K and Shivendu Shivendu (2005), "Managing Piracy: Pricing and Sampling Strategies for Digital Experience Goods in Vertically Segmented Markets," Information Systems Research 16(4), 400-417

[8] Cheng, Hsing Kenneth, Shengli Li and Yipeng Liu (2015), "Optimal Software Free Trial Strategy: Limited Version, Time-locked, or Hybrid?," Production and Operations Management 24(3), 504-517.

[9] Cheng, Hsing Kenneth and Qian Candy Tang (2010), "Free Trial or No Free Trial: Optimal Software Product Design with Network Effects," European Journal of Operational Research 205(2), 437-447.

[10] Cheng, Hsing Kenneth and Yipeng Liu (2012), "Optimal Software Free Trial Strategy: The Impact of Network Externalities and Consumer Uncertainty," Information Systems Research 23(2), 448-504.

[11] Ching, Andrew T (2010), "Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs after Patent Expiration," International Journal of Industrial Organization 28(6), 619-638.

[12] Ching, Andrew T, Tülin Erdem and Michael P. Keane (2013), "Learning Models: An Assessment of Progress, Challenges, and New Developments," Marketing Science 32(6), 913-938.

[13] Chintagunta, Pradeep K., Renna Jiang and Ginger Z. Jin (2009), "Information, Learning, and Drug Diffusion: The Case of Cox-2 Inhibitors," Quantitative Marketing and Economics 7(4), 399-443.

[14] Coscelli, Andrea and Matthew Shum (2004), "An Empirical Model of Learning and Patient Spillovers in New Drug Entry," Journal of Econometrics 122(2), 213-246.

[15] Crawford, Gregory and Matthew Shum (2005), "Uncertainty and Learning in Pharmaceutical Demand," Econometrica 73(4), 1137-1173.

[16] Dey, Debabrata, Atanu Lahiri and Dengpan Liu (2013), "Consumer Learning and Time-Locked Trials of Software Products," Journal of Management Information Systems 30(2), 239-267.

[17] Dickstein, Michael J. (2018), "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," Working Paper.

[18] Dubé, Jean-Pierre H., Günter J. Hitsch and Peter E. Rossi (2010), "State Dependence and Alternative Explanations for Consumer Inertia," The RAND Journal of Economics 41(3), 417-445.

[19] Erdem, Tülin and Michael P. Keane (1996), "Dicision Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," Marketing Science 15(1), 1-20.

[20] Erdem, Tülin, Michael P. Keane and Baohong Sun (2008), "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality," Marketing Science 27(6), 1111-1125.

[21] Fernandez-Villaverde, Jesus and Juan F. Rubio-Ramirez (2007), "Estimating Macroeconomic Models: A Likelihood Approach," Review of Economic Studies, 74(4), 1059-1087.

[22] Foubert, Bram and Els Gijsbrechts (2016), "Try It, You'll Like It - Or Will You? The Perils of Early Free-Trial Promotions for High-Tech Service Adoption," Marketing Science 35(5), 810-826.

[23] Fourt, Louis A and Joseph W. Woodlock (1960), "Early Prediction of Market Success for New Grocery Products," Journal of Marketing 25(2), 31-38.

[24] Goettler, Ronald L. and Karen Clay (2011), "Tariff Choice with Consumer Learning and Switching Costs," Journal of Marketing Research 48(4) 633-652.

[25] Grubb, Michael D. and Matthew Osborne (2015), "Cellular Service Demand: Biased Beliefs, Learning and Bill Shock," American Economics Review 105(1), 234-271.

[26] Heiman, Amir and Eitan Muller (1996), "Using Demonstration to Increase New Product Acceptance: Controlling Demostration Time," Journal of Marketing Research 33(4), 422-430.

[27] Iyengar, Raghuram, Asim Ansari and Sunil Gupta (2007), "A Model of Consumer Leaning for Service Quality and Usage," Journal of Marketing Research 44(4), 529-544.

[28] Jing, Bing (2016), "Consumer Evaluation Costs, Product Differentiation, Price Competition," Marketing Science 35(1), 113-127

[29] Johnson, Justin P. and David P. Myatt (2006), "On the Simple Economics of Advertising, Marketing and Product Design," The American Economics Review 96(3), 756-784.

[30] Keane, Michael P and Kenneth I. Wolpin (1994), "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," The Review of Economics and Statistics 76(4), 648-672.

[31] Lee, Clarence, Vineet Kumar and Sunil Gupta (2017), "Designing Freemium: Balancing Growth and Monetization Strategies," Working Paper.

[32] Lee, Robin (2013), "Vertical Integration and Exclusivity in Platform and Two-sided Markets," American Economic Review 103(7), 2960-3000.

[33] Lewis, Tracy R and David E. M. Sappington (1994), "Supplying Information to Facilitate Price Discrimination," International Economics Review 35(2), 309-327.

[34] Liu, Charles Zhechao, Yoris A. Au and Hoon Seok Choi (2014), "Effects of Freemium Strategy in the mobile App Market: An Empirical Study of Google Play," Journal of Management Information Systems 31(3), 326-354.

[35] Magnac, Thierry and David Thesmar (2002), "Identifying Dynamic Discrete Decision Processes," Econometrica, 70(2), 801-816.

[36] Nair, Harikesh (2007), "Intertemporal Price Discrimination with Forward-looking Consumers: Application ot the US Market for Console Video-games," Quantitative Marketing and Economics 5(3), 239-292

[37] Narayanan, Sridhar, Pradeep K. Chintagunta and Eugenio J. Miravete (2007), "The Role of Self Slection, Usage Uncertainty and Learning in the Demand for Local Telephone Service," Quantitative Marketing and Economics, 5(1), 1-34.

[38] Narayanan, Sridhar, Puneet Manchanda and Pradeep K. Chintagunta (2005), "Temporal Differences in the Role of Marketing Communication in New Product Categories," Journal of Marketing Research 42(3), 278-290.

[39] Niculescu, Marius F. and D.J. Wu (2014), "Economics of Free Under Perpetual Licensing: Implications for the Software Industry," Information Systems Research 25(1), 173-199.

[40] Osborne, Matthew (2011), "Consumer Learning, Switching Costs, and Heterogeneity: A Structural Examination," Quantitative Marketing and Economics 9(1), 25-70.

[41] Rekhi, Ada Chen (2017), "Freemium and Free Trial Conversion Benchmarks," Medium, Available from: https://adachen.com/freemium-and-free-trial-conversion-benchmarks-9fab38d096da

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

[43] Ryan, Stephen P and Catherine Tucker (2012), "Heterogeneity and the Dynamics of Technology Adoption," Quantitative Marketing and Economics 10(1), 63-109.

[44] Skok, David (2015), "2015 Pacific Crest SaaS Survey - Part 1," Entrepreneurs, Available from: https://www.forentrepreneurs.com/2015-saas-survey-part-1/

[45] Soysal, Gonca P, Lakshman Krishnamurthi (2012), "Demand Dynamics in the Seasonal Goods Industry: An Empirical Analysis," Marketing Science 31(2), 293-316.

[46] Stokey, Nancy L (1979), "Intertemporal Price Discrimination," The Quarterly Journal of Economics 93(3), 355-371.

[47] Turnbull, Alex (2014), "Results of The 2013 SaaS Small Business Conversion Survey!," Groove, Available from: https://www.groovehq.com/blog/saas-conversion-survey-results

[48] Wei, Xueqi and Barrie R. Nault (2013), "Experience Information Goods: "Version-to-Upgrade"," Decision Support Systems 56, 494-501.

[49] Zhu, Feng and Xiaoquan (Michael) Zhang (2010), "Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics," Journal of Marketing 74(2), 133-148.