

Service Quality Variability and Termination Behavior¹

S. Sriram (University of Michigan)
Pradeep Chintagunta (University of Chicago)
Puneet Manchanda (University of Michigan)

This version: October 2012

PRELIMINARY AND VERY INCOMPLETE

¹ We thank seminar participants at Washington University, St. Louis, Rice University, Stanford University and the University of Utah for their comments and suggestions. The second author thanks the Kilts Center for Marketing at the Booth School of Business, University of Chicago for financial support. The authors are solely responsible for all errors and omissions.

Abstract

Service quality plays a critical role in determining customer retention with good service quality typically associated with higher customer lifetime values. Given the intrinsic variability in service encounters, both the average service quality and the extent of variation around this mean are likely to be of consequence. While the effect of average service quality on customer retention has been studied, there is very little empirical research on the effect of service quality variability. Challenges to estimating the impact of service quality variability stem from not observing objective measures of service quality and behavior over time for households; and from consumers self-selecting into different quality levels. We overcome these challenges using a unique dataset that records a key dimension of quality for a video-on-demand service – the number of new movies available to a household in each time period. Importantly the number of movies available varies exogenously for a household over time. Together, these features make our data amenable for the measurement of service quality on household behavior.

Patterns in our data suggest the following. First, as expected, a higher quality level is associated with lower termination. Further, a higher variance in quality raises termination rates among households indicating that, on average, households are “risk averse,” i.e., they do not like variability in quality. Importantly, the data also show an interaction, i.e., households receiving low average quality levels with high variance terminate less than those who receive the same average quality level at a lower variance. We also find that households’ termination probabilities increase in the number of “bad” quality periods (but not with the number of “good” quality periods). Since households are required to make subscription / termination decisions for a given month prior to observing their qualities for that month, households face uncertainty about the quality they will receive that period as well as the average quality they receive as a household. The latter source of uncertainty is resolved over time via learning. Consequently, the household has an incentive to subscribe and retain the option of terminating the service in the future. The key tradeoff confronting a household with variability in quality is the following – more variability reduces utility and thereby encourages termination. On the other hand, more variability deters learning so households that have a high prior for the quality of the service are unlikely to terminate because they cannot learn about lower true quality.

Next, we investigate whether this tradeoff is manifested in our data. Identification of risk aversion is therefore important for our analysis. In our data, the ability to observe the variability in quality enables us to identify risk aversion, if any. We find that for households that receive low quality on average but with high variability, the learning deterrence effect is bigger than the risk aversion effect – higher variability leads to lower termination rates. On the other hand, when households receive a high quality level, the risk aversion effect dominates leading to increased variability increasing termination. Thus, one explanation for the interaction observed in our data is the tradeoff between learning and risk-aversion. Using the estimates from the model, we then characterize the service elasticities for termination and usage. Finally, we derive implications for the firm providing the service in terms of raising / lowering the service quality

1. INTRODUCTION

If customers of a service experience low *level* of quality, on average, it will likely result in higher termination rates because of disillusionment with the service and in some instances, due to the inability to use it (e.g., cellphone usage for a traveling sales person depends on the quality of the service in the places traveled to). This need for and importance of maintaining a high level of service quality is echoed in the academic marketing literature (see e.g., Bolton et al. 2006).

Researchers in marketing have also investigated the effects of service quality *variability* on customer retention (see Kannan and Proenca 2010 for a review of this literature). Unlike the recommendation about the level of service quality, the effect of variability on consumer retention has been viewed more ambiguously in the literature. The dominant view is that customers are likely to penalize a service for high variability, possibly because of risk aversion, resulting in lower adoption (Meyer 1981) and retention (Rust et al. 1999). On the other hand, researchers have argued that variability can increase retention if it induces a significant number of positive customer experiences (e.g., Bolton et al. 2006). Moreover, as Sun (2012) demonstrates, when customers are uncertain about the true quality of a product/service at the time of adoption, there could be an interaction effect between *variability* (in this case, of customer ratings of a movie that a potential moviegoer is exposed to) and the *mean* quality (i.e., average customer rating) on the decision. At the higher end of the average quality spectrum, high variability can lower adoption, possibly due to risk aversion. On the other hand, at lower average levels, high variability can lead to higher adoption by informing customers that it is a niche product/service. Meyer (1981) demonstrates such an interaction effect between average quality and variability although he does not find a beneficial effect of variability.

In this paper, we investigate the roles of the level and variability in service quality in driving customer retention. Apart from the main effects, we empirically document the presence of an interaction effect between average service quality and its variability on termination rates; customers

who receive a low mean quality level with high variability have a higher retention rate than those with the same quality levels but with low variability. We then postulate a mechanism involving risk aversion, learning, and consideration of option value which can induce this interaction effect. Since customer retention deals with consumers' own experiences of service quality post-adoption rather than the information contained in others' experiences pre-adoption, the "niche appeal" argument advanced by Sun (2012) does not apply in our case.

The context of our application is a video on demand (VOD) service. Customers, who are new to the service, experience it post-adoption and make monthly subscription or termination decisions. Our data reveal an interaction effect, between average service quality and its variability, on termination rates. Whereas high variability is associated with greater termination rates when average quality is high, it lowers termination amongst customers experiencing low average quality. We begin by ruling out two plausible explanations: (a) customers with high variability at low quality levels receive higher average quality than their low variance counterparts and (b) higher retention amongst high variability customers is driven by the realization of some very positive experiences (as in Bolton et al. 2006).

Our explanation for the interaction effect is as follows. First, consistent with the literature, customers are intrinsically risk averse and thus dislike variability; a feature supported by our data patterns. Further, in the context of a new service, customers tend to be uncertain about the average service quality and learn about this quality over time; we find support for this assumption in our data as well. In the presence of uncertainty and learning, high variability can deter learning about average quality. This is likely to induce a differential effect of variability based on whether the mean quality received by a household is higher or lower than their prior belief about the service. Households that have a higher prior belief about the quality of the service vis-à-vis what they actually receive are less likely to terminate in the presence of high variability. Consistent with this notion, our data show differential learning patterns among households that are subject to high vs. low variability with the

latter displaying faster learning. Thus, the key tradeoff is the following – high variability reduces the expected utility via risk aversion and thereby encourages termination. On the other hand, high variability can deter learning about the true quality of the service, which can lead to higher retention amongst customers with high prior quality beliefs who receive low quality. An added element is that in the presence of learning, consumers may be willing to stay with the service longer in order to preserve the option value of termination at a future date with more precise information about the true quality (Erdem and Keane 1996, Hitsch 2006). This is even more likely to be true if termination is an irreversible process or if reactivation involves a high transaction cost (as is true in our case). High variability can enhance the benefit from preserving option value. Together, risk aversion, learning, and consideration of option value are likely to influence the way customers react to variability in service quality.

One of the challenges in studying the impact of service quality variability is the difficulty for researchers to obtain measures of the quality of encounters and behavior over time at the individual customer level. Furthermore, even if such data are available, consumers are likely to self-select into different quality levels. We overcome these challenges using a unique dataset that records a key dimension of quality for a video-on-demand (VOD) service at the household level over time. Notably, the quality of the terrestrial signals varies exogenously across households based on their location. This enables us to understand the effect of average service quality. Furthermore, based on exogenous factors such as weather conditions, each household experiences exogenous temporal variation in signal quality. This feature of the data enables us to understand the effect of variability on household behavior. Specifically, observing the within-household temporal variation enables us to identify and estimate risk-aversion, which is key to inferring the adverse effect of variability.

Next, we calibrate a model incorporating the roles of risk aversion, learning, and the consideration of option value and demonstrate the differential effects of variability on customer retention. We calibrate both myopic and forward-looking versions of the model. While the myopic

model accounts for risk aversion and learning, the forward-looking models considers the role of option value in addition. Our results reinforce the notion that high service quality is associated with lower termination rates. Moreover, our estimates suggest that households exhibit risk aversion and learn about the quality of the service over time. Importantly, both models can replicate the interaction effect of mean and variability in signal quality on termination.

Based on the model and estimation results, we document that the in the context of new services where customers are likely to learn about their quality, households that experience low variability in service are likely to be more responsive to the quality level. It is this differential responsiveness that results in the interaction effect between service quality level and its variability (as in Meyer 1981). Furthermore, we show that this interaction effect is likely to result in lower termination rates amongst customers experiencing high variability under two conditions: (a) there exist customers for whom prior beliefs about quality are higher than the actual quality they experience and (b) the overall variability in quality is not very high. Finally, we show that increasing service quality variability does not always have to increase termination: amongst customers experiencing low variability, if the prior belief is sufficiently high, increasing variability can actually lower termination rates.

The rest of the paper is organized as follows. We first discuss the data used in our analysis and describe some key patterns. We then present the model, discuss its implications, and review the estimation approach. Next, we present our empirical results and discuss their implications. Finally, we provide some concluding comments.

2. DATA

The data used in this study come from a video on demand service which was test marketed for over a year in three cities: Jacksonville, FL, Salt Lake City, UT and Spokane, WA; and span the first 13 months of

this test market from October 2003 to October 2004.² The service allowed subscribing households to rent movies from a database of 100 titles that reside in a set-top box with a built-in antenna. During each month, the movie database was to be updated by up to a maximum of 40 new titles via a terrestrial signal from a local broadcasting tower to the set-top box. The one-time activation fee was \$29.99, the monthly subscription fee for the service was \$7.99 and the average per-movie rental fee was \$1.99. Not paying the subscription fee in a given month is viewed as the termination of the service. Reactivating requires re-payment of the activation fee. The data contain household level information on subscription and termination each month. The data used in our analysis consists of 3246 households who adopted the service of which 695 terminated it during the period of our analysis. None of the terminating households re-adopted the service. Of the households in our sample, 44.1% were from Salt Lake City, 43.3% were from Jacksonville and the remaining 12.6% were from Spokane.

A unique characteristic of the service is that while 40 movies can *potentially* be updated in its database every month, the actual number that gets updated depends on the signal quality that the household receives during that period.³ The set top box records the signal quality received by each household during each period and our data contain this information.⁴ A survey of 8% of subscribers showed that the number of movies received (a) was the most important factor affecting satisfaction with the service; and (b) explained over 50% of the variation in satisfaction across subscribers. This implies that signal quality is one proxy for the overall quality of the service. For the remainder of the paper therefore, we use the terms “quality”, “service quality” and “signal quality” interchangeably. Given that the signals are transmitted terrestrially, the average signal quality varies *across* households depending on where they are located. To illustrate this point, we present the distribution of the average

² The service was subsequently launched in a few US cities.

³ Signal quality only affects the number of movies that are updated and not their picture or sound quality. In other words, if 30 movies are updated in a given month, all 30 movies will be of the expected quality.

⁴ Signal quality information is coded on a scale of 0-2.4. The data were transmitted from the set top box to the firm in real time via the telephone line

signal quality received by households in Figure 1. The figure suggests that there is considerable heterogeneity in signal quality across households, with some not getting any *new* movies during their tenure (i.e., average signal quality of 0). Furthermore, *within* a household, the signal quality could vary over time depending on factors such as the weather, changes in the orientation of the antenna atop the set-top box, etc.; and there is significant heterogeneity amongst households on the extent to which they experience variability in signal quality. In Figure 2, we provide the distribution of within-household variability in signal quality. Therefore, there is significant heterogeneity across households in terms of the average signal quality and variability that they experience over time.

Several aspects of the data make it well suited for our research. First, given that the set top box records the signal quality received by each household during each period, we are able to observe an important component of service quality over time. This helps us understand how the mean service quality and its variability affect customer tenure. Inability to track customer experiences over time and combine them with their retention has been one of the main reasons behind limited research in this area. Second, the cross-sectional and temporal variation in signal quality is induced by factors such as household's location and weather conditions, respectively. As a result, the household cannot systematically alter the quality it receives. Further, while the firm can make investments to improve the average quality across all households, it does not do so during the period of our data. Even if the firm makes such an investment, it cannot influence the level of quality received by a *particular household* in the market. Together, these factors allow us to treat signal quality as being exogenous to the subscriber. This feature of the data allows us to measure the effects of quality without worrying about endogeneity or selection issues. Finally, the data contain information from the time of service activation for all households. This enables us to investigate how customers' beliefs about the service evolve over time as they obtain more information.

2.1. Data Patterns

A central idea in this paper is that both the mean and the variability in service quality will influence how long a customer stays with the service. In this section, we will show that our household-level data reflect both these features. First, it has been documented in the literature (see Bolton et al. 2006) that service termination (continuation) rates decrease (increase) with quality; thus a household that consistently receives poor quality will likely be disillusioned and hence, terminate the service. In our case, given service variability, this would imply that termination rates decrease with *average* quality levels. The correlation between the binary outcome of termination (1) / continuation (0) for each household and that household's average signal quality over its tenure with the service is -0.29 and is statistically significant at the 5% level. In other words, households who, on average receive higher signal quality, terminate less than those receiving lower quality. As an alternative illustration of this relationship, we provide the average termination rates for households that received high vs. low signal quality in Table 1.⁵ As expected, the table shows that households receiving a high signal quality, on average, tend to terminate (statistically) significantly less than those receiving a low signal quality.

Next, we investigate the association between termination rates for households and the extent of variability in their signal quality experienced by them. Once again, this is a cross-sectional association in which we correlate the binary household-level outcomes of termination / continuation with households' signal quality variances over their durations of service. Our data indicate a positive correlation of 0.09, which is statistically significantly different from 0 at the 5% level of significance. This shows that households who experience higher variability in their service are likely to terminate more than those with low variability. We confirm this in Table 2 in which we provide the average termination rates for households who receive above vs. below median signal quality variance. These results show that households with high signal quality variability tend to terminate more than those with low variability. This relationship is consistent with aversion to variance or more generally risk aversion.

⁵ We classified households into high and low quality groups based on a median split. We defined the median from the distribution of average signal quality across households.

To understand whether the effect of variability in service quality differs across different levels of (average) quality received, we compare the termination rates at high and low levels of both mean and variance in signal quality. The results in Table 3 suggest an interaction effect of mean and variance; whereas high variability is associated with higher termination rates at high mean quality levels, it *lowers* termination at lower quality levels. To further refine our analysis, we classified households into 5 bins based on the average quality received and compared the termination rates within each of these bins across high and low variability households. We present the results from this analysis in Figure 3. The results are consistent with the previous finding. Comparing across the high and low variability levels at high average quality levels, we find that high variability leads to higher termination rates than low variability. However, comparing across variability levels at low average quality levels, we find the opposite, i.e., higher variability leads to lower termination rates at low average quality levels. In other words, high variability can increase termination rates at high quality levels but helps in retaining customers who experience low quality, on average.

An obvious confound to our above results at low quality levels would be if households with higher variability also receive a higher signal, i.e., more movies than those who are subject to low variance. If true, it would trivially explain why termination rates are lower when variability is high. In Table 4 we provide the termination rates at different levels of average signal quality for both low variance and high variance households. By performing this analysis on a very fine grid of average quality levels, we ensure that the two groups do not differ much on this dimension. Nevertheless, the termination rates amongst households experiencing high variability is lower (higher) at low (high) average quality levels. Thus the mechanism by which higher variability lowers termination requires further investigation.

To establish the above interaction effect more formally, in Table 5 we provide the results from estimating a log-logistic hazard model to the cross-sectional data. In particular, the dependent variable

is the time to termination for each household, with those households not terminating the service by the end of the horizon surviving with the service (thus the estimation needs to account for right-censoring of the data). The key covariates included in the analysis are the average signal quality received, quality variance, and interaction between mean and variance. Further, we control for cross-sectional differences by including household demographics as well as dummy variables for the month in which the household activates the service. The results reveal that all 3 covariates have statistically significant effects on time to termination. Our results once again point to the presence of a significant interaction effect that requires explanation.

The above analysis essentially looks at cross-sectional variation across households in our sample and does not reflect how termination behavior varies based on the history of realized signal quality. If households are uncertain about the quality they are receiving, on average, the temporal variation in signal quality within a household implies that a single realization is not sufficient to fully resolve this uncertainty. Rather, each realization provides a (noisy) signal of the true average quality being received by the household. Under these circumstances, households could be learning over time about the true quality of the service they receive. In Figure 4, we plot the proportion of subscribers who terminate the service after receiving “n” periods of high quality signals or “n” periods of low quality signals for various values of “n.” Here a “high” or “low” quality signal is defined as one that falls above or below (respectively) of the median signal quality. The median is computed across all signals received by all households over the data period. The figure shows a higher termination rate for those who receive bad quality signals with the effect becoming more pronounced with the number of periods for which the quality is low. At the same time, termination rates decline slightly but do not appear to change much when households receive many periods of good signals. Together, these patterns seem to indicate that households learn about the quality they receive from the service over time.

Next, we look at how the patterns in Figure 4 vary by households subject to low and high variability. These are depicted in Figure 5. This figure suggests that households that have low variability and receive low quality signals over several time periods tend to terminate a lot more than households receiving high quality signals (i.e., there is limited overlap in their 95% confidence bands). On the other hand, there is no significant difference in termination rates based on the history of receiving low or high quality signals amongst households experiencing high variability (i.e., there is significant overlap in their 95% confidence bands). Thus, a household with high variability consistently receives several high quality signals, its termination rate is not significantly lower than one that receives several low quality signals. Together, these results seem to suggest that the rate of learning is tied to the variability in signal quality experienced by households. A plausible explanation is that while low variability enables households to discern the true quality of their signals, high variability renders such an inference difficult. In a Bayesian learning sense, this is likely to arise because the posterior will put more weight on the prior than it will on the signal.

To summarize, our investigation into the data reveal the following: higher average signal quality lowers termination rates whereas greater variability increases termination rates, i.e., consumers are risk averse. Importantly, we find that there is a significant interaction effect between mean and variance – termination rates are lower when variability is high but only at low average signal quality levels. This interaction cannot be attributed to higher variance households at low quality levels receiving more movies to watch. Our investigation into the temporal pattern of termination rates reveals that households appear to be learning about their signal qualities over time based on the realization of signals in each month; however this learning is inhibited when variability is high. In the latter case, households that have a high *prior* quality valuation of the service will likely stay with the service even when receiving low quality (on average). Thus our data reveal two key phenomena – risk aversion and learning about signal quality and the possible tension between them – at low signal quality levels, a high

variability makes one more likely to terminate due to risk aversion but could also lower termination by inhibiting learning. In the next section, we will show how these two phenomena can jointly explain the interaction effect between signal mean and variance in how they impact termination rates that we also observe in our data.

3. MODEL

3.1. Overview

We begin this section by discussing a general model of a household's subscription vs. termination decision during each period. We then lay out the specific model with risk aversion, learning, and consideration of option value, that we take to the data.

We characterize the indirect utility that household h derives from subscribing to the service at time t as

$$\tilde{V}_{1ht} = \alpha_h + g(\tilde{Q}_{ht}) + \delta D_h + \tau S_t + \varepsilon_{1ht}, \quad (1)$$

where, α_h is the household's intrinsic preference for the service, \tilde{Q}_{ht} is the household's perception about the quality of the service at time t and influences the utility from subscribing to the service via the function $g(\cdot)$, D_h are household-specific demographic characteristics that shifts their intrinsic preference for the service, and S_t are month dummies that capture seasonality in preference for the service. The term ε_{ht} is an idiosyncratic shock experienced by household h at time t .

Since the subscription decision is made at the beginning of the period, the household is uncertain about the quality of the service it will receive over the course of the month and whether it is worth the subscription fee. As discussed in the data section, there is significant heterogeneity across households in terms of the average quality of the service that they receive. Furthermore, there is temporal variation in quality experienced by a given household over time. In light of these, the household is likely to be uncertain about (a) the average quality of the service it experiences and (b) the

actual quality it will experience during period t for which it is making the subscription decision. Of these, the latter arises because of the temporal variation induced by weather conditions. In view of these two uncertainties about quality, the household's decision is based on the *expected* utility from subscribing, where the expectation is taken over the distribution of uncertainty. Formally, the expected utility that the household derives from subscribing to the service at time t is:

$$\bar{V}_{iht} = E[\tilde{V}_{iht}] = \alpha_h + E[g(\tilde{Q}_{ht})] + \delta D_h + \tau S_t + \varepsilon_{iht}. \quad (2)$$

In our empirical application, we assume that both sources of uncertainty come from a normal distribution. Specifically, we assume that at time t , the uncertainty about the household's true quality is distributed $N(\bar{Q}_{ht}, \sigma_{Qht}^2)$ and the temporal uncertainty is distributed $N(0, \sigma_h^2)$. Over time, households accumulate experience with the service, hence both the mean and the variance \bar{Q}_{ht} and σ_{Qht}^2 will evolve over time (hence the time subscript); given our assumption of Bayesian learning, the mean converges to the true average quality received by the household and the variance converges to 0. We normalize the observed (to the researcher) component of the per-period flow utility from not subscribing to the service to 0. Therefore, the total flow utility from terminating the service, including the idiosyncratic shock is

$$V_{0ht} = \varepsilon_{0ht}. \quad (3)$$

3.2. Risk Aversion

The expression for the expected utility will depend on the specification of the $g(.)$ function. If $g(.)$ is specified as linear function of \tilde{Q}_{ht} , the expected utility will be a linear function of the household's perception about the quality of the service, \bar{Q}_{ht} . Furthermore, the expected utility will be invariant to the uncertainty about service quality. As discussed in the previous section, our patterns suggest that households are averse to variability, on average. In order to capture this, we need to incorporate risk

aversion by specifying a non-linear functional form for $g(\cdot)$. In our application, we use a quadratic specification for $g(\cdot)$ such that $g(\tilde{Q}_{ht}) = \beta_h \tilde{Q}_{ht} - \gamma \tilde{Q}_{ht}^2$ where β_h captures the household's sensitivity to quality and $\gamma (> 0)$ is the coefficient of risk aversion. We use the quadratic specification because researchers have used it in the past to effectively capture risk aversion (see, for example, Erdem and Keane 1996). Furthermore, it is computationally easier to estimate a model with quadratic utility as opposed to other functional forms such as constant absolute risk aversion (CARA), especially for the model with option value.⁶ The corresponding expression for the expected utility in Equation 2 is

$$\bar{V}_{1ht} = E[\tilde{V}_{1ht}] = \alpha_h + \beta_h \bar{Q}_{ht} - \gamma \bar{Q}_{ht}^2 - \gamma \sigma_{ht}^2 + \delta D_h + \tau S_t + \varepsilon_{1ht}. \quad (4)$$

In the above expression, $\sigma_{ht}^2 = \sigma_{Qht}^2 + \sigma_h^2$ is the sum of the household's perceived uncertainty about the average quality of the service it receives (σ_{Qht}^2) and the temporal variation it experiences (σ_h^2). Together, these capture the total uncertainty perceived by the household. Recall that we previously noted that one of the contributions of the paper is the ability to estimate the risk-aversion parameter γ . In (4) a simple way to verify identification is to look at the case when $\sigma_{Qht}^2 \rightarrow 0$ over time via learning. In that case, $\sigma_{ht}^2 \rightarrow \sigma_h^2$. Since we observe the variance of quality received by individual households, σ_h^2 is data in the above equation and the variation in this variance across households allows us to identify γ . Note that we are conservatively assuming that identification comes only from cross-sectional variation in the variance across households. In reality, γ can also be identified due to changing σ_{Qht}^2 .

3.3. Learning

As discussed in the previous section, our data patterns suggest that households that consistently receive poor quality signals have a higher propensity to terminate. Furthermore, this effect is more pronounced

⁶ While not reported here (available from the authors), results from the CARA assumption are very similar to those from the quadratic model when there is no forward-looking behavior.

as they accumulate more evidence of poor service quality. Together, these imply that households might be learning about the service as they gain more experience with it.

As in the literature, we assume that households update their beliefs about the true quality of the service in a Bayesian manner. Specifically, we assume that the household has a prior belief that the quality of the service is distributed $N(Q_0, \sigma_{Qh0}^2)$. During each period t , household i notes the quality of the signals it receives (via the number of new movies in its set-top box) and updates its belief about the service. As is common in the literature (see, for example Coscelli and Shum 2004), we assume that the quality of the signals that the household receives each period comes from a normal distribution with variance σ_h^2 .⁷ A key departure from the literature, however, is that we, as researchers, observe the actual signal quality received by each household in each period.

Given the distributional assumptions, we can use Bayes rule to obtain closed form expressions for the posterior belief that a household would have about the quality of the service during each period since activation (DeGroot 1970). Especially, the conjugacy of the normally distributed prior and the normally distributed information that the household accumulates, the posterior belief is also normally distributed. Formally, the posterior mean belief about the signal quality after t periods of subscribing (i.e., the prior at the beginning of period $t+1$) based on the information it has accumulated till time t can be written as

$$\bar{Q}_{ht} = \bar{Q}_{ht-1} + \frac{\sigma_{Qh,t+1}^2}{\sigma_{Qh,t+1}^2 + \sigma_h^2} [Q_{ht} - \bar{Q}_{ht-1}] \quad (5)$$

where, Q_{ht} is the actual signal quality experienced by the household during period t and the posterior variance at the end of the period (i.e., prior variance in period $t+1$) is

$$\sigma_{Qh,t+1}^2 = \frac{1}{\frac{1}{\sigma_{Qh,t}^2} + \frac{1}{\sigma_h^2}}. \quad (6)$$

⁷ Note that we assume the variance in signals as perceived by the household is the same as temporal variation in experiences. Subsequently, we verify the sensitivity of our key results to this assumption by incorporating uncertainty in this variance.

Of note in the above expression is that the number of periods that household i is with the service does not correspond to the number of periods that the service is active unless a household activates the service in the first time period. Thus the value of t will be specific to each household.

3.4. Option Value

An implication of learning is that the perception that household h holds about the average quality of the service it receives, \bar{Q}_{ht} , will get updated over time. Moreover, the household's perceived uncertainty about the average quality of the service it receives will decline over time as it accumulates more experience. The improved precision in beliefs as it accumulates more experience implies that in addition to the contemporaneous expected utility from subscribing, the household also derives an option value from continuing with the service. The option value arises because the household can gather more information about the service by subscribing and hence make a more informed decision regarding its value in the next period (Hitsch 2006).

Given the dichotomous nature of the decision, this option value is the discounted expected value of the maximum of the two options (i.e., subscribe vs. terminate) with the information that the household expects to have in the next period. Formally, the total expected utility (including the option value) that the household would derive from subscribing can be written in the form of a Bellman equation as

$$\begin{aligned} W_h^1(\Omega_{h,t}) &= \bar{V}_{1ht} + \varepsilon_{iht} + \delta \int E_\varepsilon[\max\{W_h^1(\Omega_{h,t+1}), W_h^0(\Omega_{h,t+1})\} | \Omega_{h,t}] d(\Omega_{h,t+1} | \Omega_{h,t}) \\ &= U_{1h}(\Omega_{h,t}) + \varepsilon_{iht}, \end{aligned} \tag{7}$$

where, $\Omega_{h,t}$ captures the set of state variables that influence household h 's utility from the service at time t and δ is the discount factor. Recall that the term \bar{V}_{1ht} captures the per-period flow utility that the household expects to derive from subscribing to the service at time t .⁸ The terms $W_h^0(\cdot)$ and $W_h^1(\cdot)$ capture the total utility from terminating and subscribing to the service, respectively, as a function of

⁸ We discuss the elements of the state space in the estimation section.

the state variables and $\delta E_\varepsilon[\max\{W_h^1(\Omega_{h,t+1}), W_h^0(\Omega_{h,t+1})\}|\Omega_{h,t}]$ reflects the option value from subscribing to the service with the expectation taken over the distribution the unobserved component of the utility, ε_{iht} . The integration in Equation (7) is performed over the distribution of the state variables at time $t+1$ ($\Omega_{i,t+1}$) conditional on the observed values at time t ($\Omega_{i,t}$) and reflects the stochastic nature of their evolution. $U_{1h}(\Omega_{h,t})$ is the total observed (by the researcher) component of the utility that household h derives from subscribing to the service given the state variables, $\Omega_{h,t}$.

The household will subscribe to the service in a given period if the total expected utility from subscribing (including the option value) during that period exceeds the corresponding expected utility from not subscribing. Since we normalize the observed component of the flow utility from terminating the service is 0 (please see Equation 3), the household would subscribe if

$$W_h^1(\Omega_{i,t}) = U_{1h}(\Omega_{h,t}) + \varepsilon_{iht} > \varepsilon_{0ht}. \quad (8)$$

In our empirical application, we assume that the household-specific idiosyncratic shocks, ε , follow a type I extreme value distribution. Given the above distributional assumption and the normalization of the utility from terminating the service, Equation 7 can be rewritten as

$$\begin{aligned} W_h^1(\Omega_{h,t}) &= \bar{V}_{1ht} + \varepsilon_{iht} + \delta \int \ln[1 + \exp(U_{1h}(\Omega_{h,t+1}))] d(\Omega_{h,t+1}|\Omega_{h,t}) \\ &= U_{1h}(\Omega_{h,t}) + \varepsilon_{iht}. \end{aligned} \quad (7')$$

The corresponding probability of subscribing is:

$$Prob_{it} = \frac{\exp(U_{1h}(\Omega_{ht}))}{1 + \exp(U_{1h}(\Omega_{ht}))}. \quad (9)$$

3.5. Discussion

3.5.1. Implications for the Myopic Model

The model presented above implies that the uncertainty experienced by households vis-à-vis service quality is likely to influence the expected utility from subscription in three different ways. First, as is

evident from Equation 4, the expected per-period flow utility is a decreasing function of the total uncertainty, $\sigma_{ht}^2 = \sigma_{qht}^2 + \sigma_h^2$, because of risk aversion. Consequently, households that experience high variability in their signals (i.e., large σ_h^2) are more likely to terminate because of risk aversion compared to those with low variability. The temporal variation in signal quality has another effect via Bayesian learning. From Equation 6, we can see that the rate at which a household's uncertainty about its average quality reduces as it accumulates information is a decreasing function of σ_h^2 . Since the uncertainty about average quality (σ_{qht}^2) enters the total uncertainty, $\sigma_{ht}^2 = \sigma_{qht}^2 + \sigma_h^2$, the presence of high temporal variation in service quality is likely to have an additional adverse effect via risk aversion. Note that the posterior uncertainty about one's own average quality is bounded from above by the prior variance, σ_{qh0}^2 , i.e., $\sigma_{qht}^2 < \sigma_{qh0}^2 \forall t > 0$. Hence, if the prior variance is relatively small compared to the temporal variation, the second adverse effect of high variability has a limited impact on termination.

Second, Equation 5 suggests that the rate at which households update their beliefs about the true quality of the service will be a decreasing function of the variability in their signals (i.e., σ_h^2). Therefore, if two households have the same average quality, the household with lower σ_h^2 will converge to this quality faster than the one with higher σ_h^2 . Thus, households that experience low variability are likely to be more responsive to quality than those with high variability. This is the source of interaction effect that we documented in the data section.

Implication 1: For new products or services where households are uncertain about the quality they receive, those that experience low temporal variability in quality are likely to be more responsive (in terms of termination) to the average quality level compared to those experiencing high variability. This leads to an interaction effect between average quality and variability.

Since the households in our data voluntarily signed up for the service, they must have had a reasonably high prior expectation about the quality of the service (or a very large positive shock at the time of activating the service although this is less likely for all the subscribers). Moreover, the firm advertised the service based on the maximum number of movies that would be updated each week. Hence, we can reasonably assume that a majority of households have a prior belief that exceeds the true quality they experience.⁹ Households with higher values of σ_h^2 are likely to stay with the service longer because their learning is deterred. On the other hand, learning deterrence is likely to adversely affect customer retention amongst households receiving higher signal quality vis-à-vis their prior belief about the service at the time of activation. Consequently, learning deterrence is likely to work in the same way as risk aversion for these households.

Implication 2: The interaction effect between quality and variability will lead to lower termination rates amongst households experiencing high variability for those households whose prior beliefs about the quality of the service are higher than the quality they actually receive.

When variability increases, so would the penalty from risk aversion. At the same time, learning deterrence would imply that for households that have higher priors than their true quality, utility from staying with the service will increase. The net effect of increasing variability would depend on the relative magnitude of the two effects. At the same time, if variability is very high, the marginal benefit of learning deterrence is likely to be relatively small for both high- and low variability households. Thus, risk aversion effect is likely to dominate very the overall service variability is very high.

Implication 3: The interaction effect will lead to lower termination amongst households experiencing high variability only if the level of variability is not too high; if variability is very high, risk aversion will dominate and high variability will always lead to higher termination even if the interaction effect exists.

3.5.2. Additional Implications for the Model with Option Value

⁹ We assess the validity of this assumption in a later section.

The impact of variability on option value is a bit more complicated. The option value is typically an increasing function of the per-period flow utility. The premise is that a household will see more future potential from a service if it yields higher flow utility and hence holds more promise. The extent to which flow utility is an increasing or decreasing function of variability is determined by the tradeoff between the adverse effect via risk aversion and the positive effect of deterring learning, at least for some households. The other driver of option value is the extent to which staying with the service helps the household to make a more informed decision in the future. If the household already has precise knowledge about the average quality it experiences, this benefit is likely to be small. Consequently, households that start with higher uncertainty about the average quality of the service and see smaller gains in precision as they accumulate more information are likely to benefit more by staying longer with the service. Thus, one can conjecture that for a given flow utility, high variability is likely to induce customers to stay longer with the service.

To summarize, there are two competing effects of variability in service quality on the per-period flow utility. On the one hand, high variability can lower the per-period flow utility via risk aversion. Furthermore, the learning deterrence effect of high variability can adversely affect the per-period flow utilities of households whose prior beliefs about the average quality that they receive are lower than what they actually realize. On the other hand, amongst households who receive low average quality, the learning deterrence effect of high variability can keep the per-period flow utility high and prolong customer lifetimes. Moreover, for a given flow utility, the presence of high variability can provide incentive for forward-looking households to stay with the service and make a more informed decision in the future. Thus, contrary to popular belief that variability is always detrimental to customer retention, our model provides a nuanced characterization of its effects.

3.6. Estimation

A typical learning model involves four parameters: (i) mean of the prior belief about the unknown entity (Q_0), (ii) prior variance (e.g., $\sigma_{Q_{h0}}^2$), (iii) variance of the signals about the unknown entity (e.g., σ_h^2), and (iv) true value of the unknown entity (Q_h). Of these, the true value of the unknown entity is estimated when the researcher does not observe the exact signals received by the households.¹⁰ Since we observe the signal strength received by individual households during each period, we do not estimate Q_h . As discussed above, since the households in our data voluntarily signed up for the service, they must have had a reasonably high prior expectation about the quality of the service. Moreover, the firm advertised the service based on the maximum number of movies that would be updated each week. In view of these we assume that the households activate the service with the belief that the maximum number of new movies that were promised in the advertisements will be available to them every week. In our data, the signal quality levels are recorded on a 0 – 2.4 scale with 0 corresponding to no new movies downloaded and 2.4 equivalent to the maximum number of movies.¹¹

We fix the signal variance, σ_h^2 , based on the actual variance in signal quality received by the household. While this might be a strong assumption, it can be justified by noting that the household can see the updates every week while our time period is a month. Consequently, by the first month, the household could technically have inferred the temporal variability it experiences and used this information to update beliefs. While this assumption is consistent with that made in the majority of the literature on learning, it is nevertheless strong for our situation. Subsequently, we verify the sensitivity of our key results to this assumption by incorporating uncertainty around this variance. Thus, the only component of the learning model we estimate is the heterogeneous prior variance about the signal quality. As discussed earlier, the ratio of the signal variance to the prior variance will help us infer the rate of learning; a higher ratio is synonymous with faster learning.

¹⁰ In these situations, the information is assumed to come from a distribution with unknown mean and variance. The researcher would then simulate the information signals by making draws from this distribution, with the unknown moments of the distribution being estimated from the data.

¹¹ In the next section, we justify the choice of this prior based on model fit.

In addition to the prior variance, we estimate the following parameters: the intrinsic preference for the service, α_h , effect of perception about signal quality, β_h , the effect of demographic characteristics, δ , eleven month fixed effects, τ , and more importantly, the risk aversion parameter, γ . In our application, we use city dummies, distance to the nearest retail store, income, number of children, and the number of elderly people in the household as the demographic characteristics that shift the intrinsic preference for the service. In order to reduce the computational burden, we discretized the continuous demographic variables by performing a median split with the high values coded as 1 and the low values coded as 0.

From Equation 4, we can see that we allow for heterogeneity in the prior variance, choice intercept, and the signal quality effect. We model this unobserved heterogeneity using the latent class specification with consumers belonging to different segments. Specifically, we assume there are R segments of households such that the parameter vector, $\Theta^r = \{\sigma_{Qr0}^2, \alpha_r, \beta_r\}$ includes the set of heterogeneous parameters for segment r , $r = 1, 2, \dots, R$. Based on Equation 9, the likelihood of observing the history of household h belonging to segment r can be written as

$$L_{hr} = \prod_{t=1}^{T_h} (Prob_{hrt})^{I_{ht}} (1 - Prob_{hrt})^{(1-I_{ht})}, \quad (10)$$

where, I_{ht} is an indicator for whether the household subscribes to the service (1) or terminates (0). The corresponding overall likelihood for all households is:

$$L = \prod_{h=1}^N (\sum_{r=1}^R \pi^r L_{hr}). \quad (11)$$

In the above expression, π_r is the size of segment r , $r = 1, 2, \dots, R$ and L_{hr} is the likelihood for household h conditional on it belonging to segment r .

3.6.1. Estimation of the Model with Option Value

The estimation of the model with option value needs elaboration. Given our model specification, there are four sets of state variables: (a) the mean belief about the quality of the service at any period, \bar{Q}_{ht} ,

(b) uncertainty about this belief $\sigma_{Q_{ht}}^2$, (c) seasonality, and (d) demographic characteristics that act as shifters of the intrinsic preference for the service. Of these, the demographic characteristics are not dynamic but require us to compute the value function for each combination of values that they can take. Given that we dichotomize the five demographic variables based on median split, our estimation involves computing 32 value functions corresponding to all possible combinations. As noted in Equation 1, the seasonality variable captures shifts in the intrinsic propensity to subscribe to the service during different months of the year. Consequently, while this state variable is intrinsically dynamic, its evolution is deterministic.

As in any learning model with option value, the two key state variables in our context are the mean and uncertainty about the unknown entity, the quality of the service. Equation 5 governs the law of motion for the belief about mean quality, \bar{Q}_{ht} . Given the distributional assumptions of normal prior and signals, as in Hitsch (2006), we can characterize the transition density of $\bar{Q}_{ht}|\bar{Q}_{ht-1}$ as $N(\bar{Q}_{ht-1}, \frac{\sigma_{\bar{Q}_{h,t+1}}^4}{\sigma_{\bar{Q}_{h,t+1}}^2 + \sigma_h^2})$. Equation 6 characterizes the evolution of the uncertainty about the average quality of the service. Unlike the mean, the uncertainty evolves in a deterministic manner as a function of the signal variance, σ_h^2 . Furthermore, as discussed elsewhere, this uncertainty always declines as a household accumulates more information.

We use the nested fixed-point approach (Rust 1987) to compute the value functions. Together, we have two deterministic state variables (demographic characteristics and seasonality) and two continuous ones. To accommodate the continuous state variables, we evaluate the value functions on a discrete grid of 20 points defined by Chebyshev zeros. We used Gaussian quadrature to evaluate the integral in Equation 7. Note that the integration requires us to evaluate the value function outside the set of grid points. To this end, we approximate the value function using a Chebyshev tensor polynomial

of degree 4. Similarly, we use Chebyshev interpolation to translate the value functions evaluated at grid points to the actual values of the state variables for individual households over time.

5. RESULTS

Table 1: Termination rates for households whose average qualities are above and below median across households

Termination rate*	
High Mean Signal	11.9%
Low Mean Signal	32.7%

Table 2: Termination rates for households whose variances in qualities are above and below median across households

Termination rate*	
Low Variance Signal	19.9%
High Variance Signal	24.7%

Table 3: Interaction effect between mean and variance

Termination rate*		
	High Signal Variance	Low Signal Variance
High Mean Signal	15.4%	10.5%
Low Mean Signal	28.5%	42.9%

Table 4: Termination Rates at Various Levels of Average Quality

LOW VARIABILITY HOUSEHOLDS			HIGH VARIABILITY HOUSEHOLDS		
AVERAGE SIGNAL QUALITY	Number of Households	% TERMINATING	AVERAGE SIGNAL QUALITY	Number of Households	% TERMINATING
0.10	50	46%			
0.24	68	47%	0.27	12	33%
0.41	65	58%	0.41	75	40%
0.58	40	70%	0.58	114	34%
0.74	36	61%	0.76	144	33%
0.92	45	69%	0.93	186	33%
1.09	42	50%	1.09	209	27%
1.26	58	14%	1.24	208	24%
1.42	104	16%	1.42	220	17%
1.59	157	13%	1.58	237	13%
1.76	222	11%	1.74	153	15%
1.92	264	14%	1.91	84	15%
2.08	256	9%	2.05	17	24%
2.25	214	14%			
2.36	38	11%			

Table 5: Results from a termination hazard model with a log-logistic baseline hazard

Dependent variable: Time to termination / end		
Variable	Coefficient	t-value
Mean signal quality	-0.62	-15.80
Signal quality variance	-0.28	-5.28
Interaction effect	0.19	3.51
Demographics	Yes	
Start period dummy	Yes	

Figure 1: Distribution of Household's Average Signal Quality across Households



Figure 2: Distribution of within household variability in signal quality




Figure 3: Termination rates for low & high variance recipients at different signal quality levels

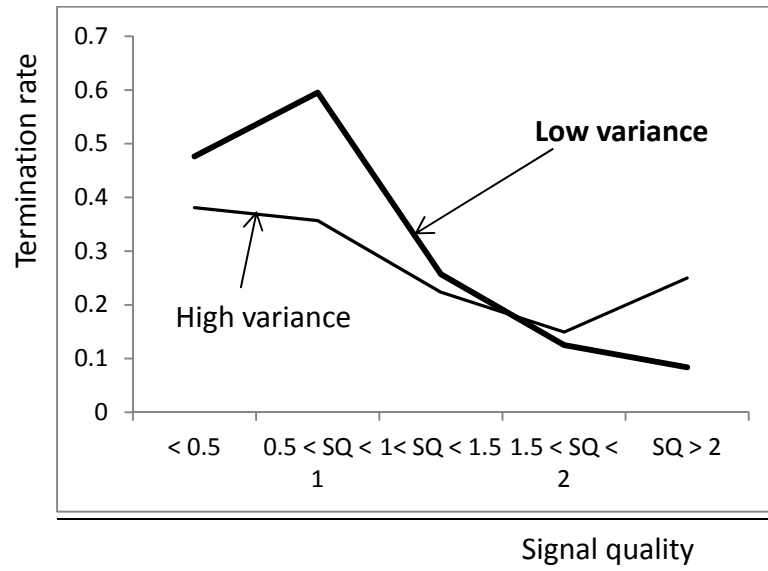


Figure 4: Evidence of Learning




Figure 5: Evidence of differential learning across low and high variability groups

