Investigating Customer Behaviour Heterogeneity Pattern Using Exponential-Gamma Timing Model

[1]akomolafe A. A. and [2]olawale A.O

[1]Department of Statistics, Federal University of Technology, Akure, Ondo state, Nigeria [akomolafe01@yahoo.com]
[2] Department of Statistics, Federal Polytechnic, Ile-Oluji, Ondo State [oludayobola@gmail.com]

Abstract: Managers today are very interested in knowing the future purchasing pattern of their customers thereby providing little insight about individual-level shopping behaviour. Additionally, behaviour may evolve over time, especially in a changing environment. This research develops an individual-level model for customer store visiting behaviour based on buying-behaviour data from standard wholesales outlet. We capture cross-sectional variation in store-visit behaviour as well as changes over time as visitors gain experience with the store, as the composition of the customer population changes, the overall degree of buyer heterogeneity that each store faces may change. We also examine the relationship between visiting frequency and purchasing propensity. However, we also show that changes (i.e., evolution) in an individual's visit frequency over time provides further information regarding which customer segments are more likely to buy. Rather than simply targeting all frequent shoppers, our results suggest that a more refined segmentation approach that incorporates how much an individual's behaviour is changing could more efficiently identify a profitable target segment.

[akomolafe A. A. and olawale A.O. **Investigating Customer Behaviour Heterogeneity Pattern Using Exponential-Gamma Timing Model.** *Researcher* 2018;10(6):61-68]. ISSN 1553-9865 (print); ISSN 2163-8950 (online). http://www.sciencepub.net/researcher. 8. doi:10.7537/marsrsj100618.08.

Keywords: Consumer Packaged Goods (CPG) data, Buying Behaviour, Duration models, Heterogeneity, Nonstationarity

Introduction

Researchers in marketing have used several different mechanisms to introduce these time-varying effects into the traditional stochastic modeling framework. For instance, Sabavala and Morrison (1981) incorporated nonstationarity by introducing a renewal process into a probability mixture model in accordance with the "dynamic inference" framework first set out by Howard (1965). Sabavala and Morrison applied this model to explain patterns of advertising media exposure over time; further applications of a similar type of renewal-process approach can be seen in Fader and Lattin (1993) as well as Fader and Hardie (1999).

Specifically, our behavioural assumption is that customers' underlying rates of buying are continually and incrementally changing from one visit to the next. As individuals adapt to and gain experience with the new retail environment, they may return to the store at a more frequent rate, a less frequent rate, or perhaps at the same rate for the next buying. By assuming that each individual will update her latent rate, λ_i, after each buying, a very simple way to specify this updating process is as follows:

$$\lambda_{i(j+1)} = \lambda_{ij}.c \tag{1}$$

Where λ_{ij} is the rate associated with individual *i*'s *j*th repeat visit and ***c*** is a multiplier that will update this rate from one visit to the next. If the updating multiplier, ***c*,** equals one, visiting rates are considered unchanging, and the stationary exponential-gamma would remain in effect. But if ***c***is greater than one, shoppers are visiting more frequently as they gain experience, and if ***c*** is less than one, shoppers are visiting less frequently as they gain experience.

However, using a constant multiplier to update the individual λ's would be a very restrictive (and highly unrealistic) way of modeling evolutionary behaviour in a heterogeneous environment. A more general approach is to replace the scalar multiplier, ***c***, with a random variable *cij* in order to acknowledge that these updates can vary over time and across people. Each individual visit will lead to an update that may increase, decrease, or retain the previous rate of visit, depending on the stochastic nature of the updating multiplier.

To generalize (1) in this manner, we assume that these probabilistic multipliers, *cij*, arise from a gamma distribution, common across individuals and visits, with shape parameter *s* and scale parameter β. We choose the gamma distribution to describe the updating multiplier for the same reasons why we used it to describe the heterogeneity in λ. It is a very flexible distribution that accommodates a variety of shapes. This gamma distribution essentially describes the nature of the behavioural evolution faced by a given store. The updated $\lambda_{i(j+1)}$ then becomes a product of two independent gamma-distributed random variables: the previous rate, λ *ij*, and the multiplier, *cij*. The overall model, therefore, uses four parameters to simultaneously capture cross-sectional heterogeneity and evolving buying behaviour: two parameters (*r* and α) govern the gamma distribution that describes the initial heterogeneity in visiting rates, and another two parameters (*s* and β).

Regardless of whether the multiplier is increasing (*cij*>1) or decreasing (*cij*< 1) a particular buying rate at a particular point in time, we expect that an individual's value will evolve relatively slowly over time. This suggests that the updating gamma distribution, *u* (*cij*; *s*, β), should have a mean fairly close to 1.0 but should also allow for more extreme increases or decreases in λ at any given update opportunity. The spread of this updating distribution is directly tied to the magnitude of the *s* and β parameters. As both of these parameters become large, the distribution degenerates towards a spike located at s/β . Taken to the extreme (i.e., *s* and β get extremely large), this model would then collapse into the deterministic updating model (1) with c= s/ β .

Finally, another interesting characteristic of the updating distribution is that it allows for customer attrition pattern, since the gamma distribution can yield a draw of *cij* extremely close to 0. When this situation arises, the customer effectively drops out and is unlikely to return to the store. Such attrition may be very common for websites and has been the centerpiece of other types of models in this general methodological area (Reinartz and Kumar 2000; Schmittlein, Morrison, and Colombo 1989). The fact that we can accommodate attrition in such a simple, natural manner is an appealing aspect of the proposed modeling approach.

3. Likelihood Specification

When estimating the ordinary (stationary) exponential-gamma model, there are two ways of obtaining the likelihood function for a given individual. The usual approach is to specify the individual-level likelihood function, conditional on that person's (unobserved) value of λ *i*. This likelihood is the product of *Ji* exponential timing terms, where *Ji* is the number of repeat buying made by panelist *i*, plus an additional term to account for the right-censoring that occurs between that customer's last arrival and the end of the observed calibration period (at time *T*):

$$L_i/\lambda_i = \lambda_i e^{-\lambda_i(t_{i1}-t_{10})}.\lambda_i e^{-\lambda_i(t_{12}-t_{i1})}\ldots\lambda_i e^{-\lambda_i\left(t_iJ_i-t_{i(J_i-1)}\right)}.e^{-\lambda_i(T-t_iJ_i)} \quad (2)$$

To get the unconditional likelihood we then integrate across all possible values of λ , using the gamma distribution as a weighting function:

$$L_i\left|r,\alpha\right. = \int_0^\infty L_i\left|\lambda_i\right..gamma\left(\lambda_i;r,\alpha\right)d\lambda_i \quad (3)$$

Where gamma (λ_i ; *r*, α) denotes the gamma distribution as shown in (1). This yields the usual exponential-gamma likelihood, which can be multiplied across the *N* panelists to get the overall likelihood for parameter estimation purposes:

$$L = \prod_{i=1}^{N}\frac{\Gamma(r+J_i)}{\Gamma(r)}\left(\frac{\alpha}{\alpha+T-t_{i0}}\right)^r\left(\frac{1}{\alpha+T-t_{i0}}\right)^{J_i} \quad (4)$$

An alternative path that leads to the same result is to perform the gamma integration separately for each of the *Ji*+1 exponential terms, and then multiply them together at the end. This involves the use of Bayes Theorem to refine our "guess" about each individual's value of λ *I* after each arrival occurs. Specifically, it is easy to show that if someone's first repeat visit occurs at time *tij*, then:

$$g\left(\lambda_{i2}\left|\text{arrival at } t_{i1}\right.\right) = gamma\left(r+1,\alpha+t_{i1}-t_{i0}\right) \quad (5)$$

The gamma distribution governing the rate of buying for subsequent arrivals follows:

$$g\left(\lambda_{i(j+1)} \middle| \text{arrival at } t_{ij}\right) = gamma\left(r + j, \alpha + t_{ij} - t_{i0}\right) \tag{6}$$

Using this logic, we can re-express the likelihood as the product of separate EG terms

$$L = \prod_{i=1}^{N} \prod_{j=1}^{J} \left(\frac{r + j + 1}{\alpha + t_{i(j-1)} - t_{i0}} \right) \left(\frac{\alpha + t_{i(j-1)} - t_{i0}}{\alpha + t_{ij} - t_{i0}} \right)^{r+j} .S\left(T - t_{iJ_i}\right) \tag{7}$$

which collapses into the same expression as (4).

When we introduce the nonstationary updating distribution, the multipliers (*cij*) change the value of λ *I* from visit to visit, thereby requiring us to use the sequential approach given in (7) to derive the complete likelihood function. We need to capture two forms of updating after each visit: one due to the usual Bayesian updating process (which is associated with stationary behaviour given by (6)) and the other due to the effects of the stochastic evolution process. Therefore, the distribution of buying rates at each repeat visit level is the product of two gamma distributed random variables – one associated with the updating multiplier and one capturing the previous visiting rate. For the case of panelist *i* making her *j*th repeat visit at time *tij*:

$$G\left(\lambda_{i(j+1)} \middle| \text{arrival at } t_{ij}\right) = gamma\left(r + j, \alpha + t_{ij} - t_{i0}\right).\, gamma\left(s, \beta\right) \tag{8}$$

One issue with this approach is that the product of two gamma random variables does not lend itself to a tractable analytic solution. However, there is an established result (see, e.g., Kendall and Stuart 1977, p. 248) suggesting that the product of two gamma distributed random variables can be approximated by yet another gamma distribution, obtained by multiplying the first two moments about the origins of the original distributions:

$$m_1^{(\lambda_{i(j+1)})} = m_1^{(\lambda_{ij})} \text{ x } m_1^{(c_{ij})}$$

and

$$m_2^{(\lambda_{i(j+1)})} = m_2^{(\lambda_{ij})} \text{ x } m_2^{(c_{ij})} \tag{9}$$

As shown in Appendix A, this moment-matching approximation, used in conjunction with Bayesian updating, allows us to recover the updated gamma parameters that determine the rate of buying, λ *ij*, for panelist *i*'s *j*th repeat visit as follows:

$$r(i, j+1) = \frac{[r(i,j)+1].s}{[r(i,j)+2].(s+1) - [r(i,j)+1].s} \tag{10}$$

$$\alpha(i, j+1) = \frac{\lfloor \alpha(i,j) + t_{ij} - t_{i(j-1)} \rfloor.\beta}{[r(i,j)+2].(s+1) - [r(i,j)+1].s} \tag{11}$$

Where *r* (*i*, 1) and s (*i,* 1) are equal to the initial values of *r* and s as estimated by maximizing the likelihood function specified in (6).

We performed 20 separate simulations to verify the accuracy of using such a moment-matching approximation. In each simulation, we first generated 1000 random draws from a gamma distribution with randomly determined shape and scale parameters to represent initial λ values. Then, a matrix of updating multipliers was also simulated for a series of five updates (i.e., five future repeat visits). Each 1000x5 matrix was generated by taking draws from a gamma distribution, again with randomly determined shape and scale parameters, where columns one through five represented the updates after one to five visits. The updated λ series after five repeat visits was calculated using two methods (1) direct (numerical) multiplication of the 1000 initial λ's and the five updating series or (2) randomly drawing 1000 values from the distribution resulting from the moment-matching approximation across all five updates. A Kolmogorov-Smirnov test of fit indicated that, for each of the 20 simulations, the distribution of values resulting from the moment-matching approximation is not significantly different from that resulting from the direct multiplication of these random variables. Therefore, we are confident that the moment-matching approximation accurately captures the gamma distributed updating process we wish to model.

After incorporating the evolution process into our model, the likelihood function to be maximized follows:

$$L = \prod_{i=1}^{N} \prod_{j=1}^{J} \left(\frac{r(i,j)}{\alpha(i,j)} \right) \left(\frac{\alpha(i,j)}{\alpha(i,j) + t_{ij} - t_{i(j-1)}} \right)^{r(i,j)+1} . S\left(T - t_{iJ_i}\right) \tag{12}$$

Where $r\ (i, j)$ and $\alpha\ (i, j)$ are defined in equations (10) and (11) while the survival function, $S\ (T\text{-}tij)$, is defined as:

$$S\left(T - t_{ij}\right) = \left(\frac{\alpha(i, J_i + 1)}{\alpha(i, J_i + 1) + T - t_{iJ_i}} \right)^{r(i, J_i + 1)} \tag{13}$$

For the special case in which behaviour is not evolving and the nonstationary updating distribution degenerates to a spike at 1.0

(i.e., $s = \beta$ = M, where M approaches infinity), then this equation collapses down exactly to the ordinary (stationary) exponential-gamma model.

4. Data

We apply the models described in the previous section to fruit juice data collected from a standard Departmental store retail outlet in Ibadan. For our purposes, we are interested in the dates of the visits each panelist makes to a given store. To consolidate the data just a bit, we aggregated buying behaviour to weekly level. For example, a shopper may leave a store briefly and return later that week. However, this second visit is unlikely to be considered a repeat visit but rather an extension of the first visit. Therefore, if a given panelist were to visit a particular store multiple times in a single week, we would encode that behaviour as just one visit for the day when the session began. Since we are interested in the timing and frequency of repeat buying, our dataset describes each customer as a sequence of weeks when buyers were made. All customers that have visited the store of interest at least once during the observation period were included in this dataset. We use data from two stores – selling juice drink and packed chicken. Fruit juice attracted 2,140 unique buyers during this six-month period totaling 9,260 visits, while packed chicken had 975 visitors making 3,210 visits (refer back to Figure 1 and Table 1 for more detailed summaries of the data).

5. Model Results

Before estimating the evolving visit model developed in §3, we first examine the static exponential-gamma timing model as a benchmark. When the static, two-parameter model is applied to the eight months of fruit juice buying behaviour data, we find that the mean rate of visit (E [λ]=r/α) is 0.0112. In other words, the expected intervisit time (1/λ) is 89.3 days, which is high, but reasonably consistent with the summary statistics mentioned earlier. But beyond their ability to capture the mean of the heterogeneous visiting process, the model parameters also provide useful information about the nature of the distribution of visit rates across the population. With a shape parameter of 0.271 and a scale parameter of 42.955, the distribution of visiting rates can be described by the gamma distribution in Figure 3. This distribution has a large proportion of the population with very low rates of buying. The median rate, according to this model, is 0.005, corresponding to an buying rate- time of 200 days. This distribution of rates is very consistent with the observed histogram of visit frequency (Figure 1), suggesting that the stationary EG model provides a very good benchmark model for visit behaviour.

A principal reason for these high expected buying rate is the fact that the stationary model does not allow shoppers to drop out and never return. As a result, a customer who has actually dropped out would be seen by the model as still being “alive,” but having a very slow visiting rate, since she would not have yet returned to the store by the end of the observation period.

In Table 2, we contrast the parameter estimates and fit statistics for the static EG model with those from our four-parameter model of evolving visiting behaviour. Not only does the latter model fit the data better, but it also has more intuitively appealing results. While the basic shape of the gamma distribution for initial buying rates (shown in Figure 4a) may appear to be similar to that of the static EG model, it is less dominated by low-frequency shoppers, leading to a substantially lower mean buying rate (52 days, E [λ] = 0.019). Likewise, the median buying rate shrinks to 167 days (median λ

= 0.006). These differences reflect the fact that dropout – or other types of evolution – can take place as the customer becomes more familiar with the product.

Table 2. Model Results for Amazon

	Stationary EG Model	Evolving Visit Model
R	0.483	0.324
α	42.955	16.857
S		2.299
β		2.304
LL	-34,347.2	-33,648.0
No. of parameters	2	4
CAIC	68,711.17	67,296.0

According to the evolutionary model, the mean update for any given visit (s/β) is very close to one (0.998) suggesting, perhaps, that it is a fairly stationary process. However, a closer look at the distribution (see Figure 4b) shows that there is significant variance about this mean. Though the mean update is close to one, the distribution is quite skewed. With a median value of c_{ij}=0.858, buyers tend to return to the store at slower rates from visit to visit. The implications of these results are in stark contrast to the measures summarized in Table 1 that implied increased visiting frequency over time.

Though other models have acknowledged the issue of nonstationarity, many of them have focused primarily on dropout (Eskin 1973, Kalwani and Silk 1980, Schmittlein, Morrison, and Colombo 1989). These models allow for individuals to make several purchases, become disenchanted, and never purchase again. To test if the evolving visit model is capturing evolving behaviour over time in addition to a dropout phenomenon, we also estimated an exponential-gamma model with a dropout component similar to that specified by Eskin (1973) and Fader and Hardie (1999).

In the EG model with dropout, the probability of buying given that you are an active buyer is modeled as an exponential-gamma process. However, the probability of being an active buyer after the *j*th visit, π_j, is determined by the following:

$$\pi_j = \Phi\left(1 - e^{-\theta j}\right) \qquad (14)$$

where ϕ is the long run probability of a customer remaining active, and θ is the rate at which the π approaches this long run probability. Though the EG model with dropout provides a significant improvement in fit over the stationary EG model (LL = -33,804.7), it does not approach the performance of the evolving visit model which has the same number of parameters.

This suggests that the evolving visit model is capturing a phenomenon in addition to just dropout.

Validation

While we have discussed the fact that the evolving visit model fares well on a relative basis compared with various benchmark models, we have yet to show that it performs sufficiently well on an absolute basis. In this section, we validate the evolving visit model by examining the accuracy of longitudinal forecasts. Because the evolving visit model relies on an approximation (9) to specify and estimate the model, we need to perform simulations to generate data for tracking/forecasting purposes. This is a straightforward and computationally efficient task. For each iteration of the simulation, we create a simulated panel that matches the actual panel in terms of its size and the distribution of its initial visit times. We then generate a sequence of repeat visits using the parameter estimates from the model. This requires us to maintain a time-varying vector of λ's for each panelist, which starts with random draws from the initial (r, α) gamma distribution, and then gets updated using the (s, β) gamma distribution after each simulated exponential arrival occurs. We continue this process until every simulated panelist gets past the tracking/forecasting horizon of interest to us. It is then a simple matter to count up the number of visits on a week-by-week basis for each iteration of the simulation. We then average across 1000 iterations to generate the tracking and forecasting plots. Using the MATLAB programming language, each of these iterations takes only a few seconds on a standard PC, and we see very consistent convergence properties after a few dozen iterations.

Before creating the forecasts, we re-estimate both models (stationary and evolving EG) using only the first half (i.e., four months) of the dataset. (It is worth noting that the evolving model parameters are quite robust to this changing calibration period, while the stationary model has a noticeably higher visit rate over the shorter period – clear evidence of the slowdown discussed earlier). To generate the forecasts for the evolving visit model, we use the simulation procedure described above. For the stationary EG model, the expected number of repeat visit per week can be calculated directly as follows:

$$E[repeat \quad \text{visits}_{\ w}] = N_w t\left(\frac{r}{\alpha}\right) \qquad (15)$$

where *Nw* is the number of eligible repeat buying in week *w* and *t* is the time period of interest, i.e., seven days in this case. Figure 5 shows cumulative forecasts as well as actual visits for the juice drink buying behaviour.

Both models seem to track the data quite well over the initial four-month calibration period. However, as we enter the forecasting period, the stationary EG model begins to diverge, ultimately over predicting by 37% for juice drink buying behaviour at the end of the eight month period. It overestimates the number of visits per week as it does not recognize that shoppers are returning less frequently over time. The evolving visit model, however, forecasts quite accurately, well within 5% of the actual sales line throughout the forecast period. This is an impressive achievement and serves as a strong testimonial to the validity of the assumptions, structure, and parameter estimates associated with the proposed model.

Results for data from PACKED CHICKEN buying behaviour

The same set of models and analyses were also applied to packed chicken buying behaviour data (results in Table 3). We see a remarkably similar set of patterns as in the case of juice drink. In moving from the static EG model to the evolving specification, we see significantly shorter intervisit times, since the latter model can accommodate customer dropout. We also see, once again, that the mean update is close to 1.0 (0.991), but with a median of 0.837, customer shopping frequency is more likely to decrease than increase after each visit. We emphasize once more that these results contradict the summary statistics from Table 1, which seemed to imply that shopping frequency is increasing from one visit cycle to the next.

Table 3. Model Results for packed chicken buying-behaviour data

	Stationary EG Model	Evolving Visit Model
r	0.255	0.165
α	28.305	8.889
s		2.084
β		2.104
LL	-9,459.6	-9,120.7
No. of parameters	2	4
CAIC	18,934.1	18,271.0

Other benchmark models (involving dropout and/or constant updates) proved once again to be vastly inferior to the evolving visit model. Finally, our forecast validation led to encouraging results with projected visits only 2% above the actual number at the end of the eight month period, compared to a 40% over-forecast for the stationary model. While we are very encouraged by these strong initial results, we are also surprised at the degree of similarity seen for these two Consumer Packaged Goods. We certainly do not want to suggest that the specific patterns captured here will generalize to all online (or offline) retailers, but this should be ample motivation for future studies to find and describe a broader range of buying-behaviour.

6. Discussion and Conclusion

Many skeptics claim that the offline or online shopping is nothing more than a new distribution channel, and thus it should not change the way we examine customer behaviour. While this may be true in certain respects, this research highlights some of the uniquely different research perspectives that we gain from examining Packaged Goods buying-behaviour data. The data available to us make it possible to study the evolution of visit behaviour. The model developed here is tailored specifically to both online and offline buying. We posit a behaviourally plausible – and highly parsimonious – model that allows visiting behaviour to evolve gradually over time, although it also allows for more abrupt changes, such as permanent dropout from the store. And indeed, our empirical analysis reveals the fact that the average update in household visiting rates is a multiplier close to 1.0, but there is significant spread around this value.

Additionally, the manner in which we implement this updating scheme – a gamma distribution to capture the different values of these multipliers – is a new methodological contribution, which merits consideration for other types of non-stationary modeling contexts.

Use of the model reveals that individual-level behaviour patterns appear to contradict the perspective that one would obtain from examining the aggregate data alone. Specifically, the aggregate data seem to indicate an acceleration of visiting behaviour at each of two leading consumer packaged goods products, yet our model parameters suggest that the typical shopper is experiencing a gradual slowdown in her buying rate over time. The difference here is that an increasing number of new visitors are coming to each store over time, masking the slowdown that may be occurring for many experienced visitors. This effect could have dramatic implications for managers who neglect to examine their data at a sufficiently fine level of disaggregation.

Beyond the intuitive appeal of the model specification and its estimated parameters, we also show that it has excellent validity from an out-of-sample forecasting perspective. For the retail stores, the model tracks future visiting patterns extremely well, remaining within 5% of the actual data over the entire duration of a four-month holdout period. While this model was not constructed with forecasting in mind as a principal objective, this result certainly speaks well about its overall versatility.

References

1. Akomolafe A.A and Amahia G.N., (2009) “customer Behaviour Transaction Data Analysis “international Journal of Management. Vol 1 (1) 32-37.
2. Angelis R. De, Capocaccia R., Hakulinen T., Soderman B. and Verdecchia A., Mixture Models for Cancer Survival Analysis: Application to Population-Based Data With Covariates, Statistics in Medicine, 18, 441-454, 1999.
3. Akomolafe A.A. (2011). Analysis of Consumer Depth of Repeat Purchasing Pattern: An Exploratory Study of Beverages Buying Behaviour Data. Journal of Business and Organisational Development.Vol 3 (1) 1-8.
4. Bellinger, Danny N., Dan H. Robertson, and Elizabeth C. Hirschman (1978), “Impulse Buying Varies by Product,” Journal of Advertising Research, 18 (December), 15-18.
5. Bronnenberg, Bart J., Vijay Mahajan, and Wilfried R. Vanhonacker (2000), “The Emergence of Market Structure in New Repeat-Purchase Categories: The Interplay of Market Share and Retailer Distribution,” Journal of Marketing Research, 37 (February), 16-31.
6. Eskin, Gerald J. (1973), “Dynamic Forecasts of New Product Demand Using a Depth of Repeat Model,” Journal of Marketing Research, 10 (May), 115-129.
7. Fader, Peter S. and Bruce G.S. Hardie (1999), "Investigating the Properties of the Eskin/Kalwani & Silk Model of Repeat Buying for New Products," in Lutz Hildebrandt, Dirk Annacker, and Daniel Klapper (eds.), Marketing and Competition in the Information Age, Proceedings of the 28th EMAC Conference, May 11-14, Berlin: Humboldt University.
8. Fader, Peter S. and James M Lattin (1993), “Accounting for Heterogeneity and Nonstationarity in a Cross-Sectional Model of Consumer Purchase Behaviour, Marketing Science, 12 (3), 304-317.
9. Howard, R.A. (1965), “Dynamic Inference,” Operations Research, Vol. 13 (2), 712-733.
10. Jarboe, Glen R. and Carl D. McDaniel. (1987). "A Profile of Browsers in Regional Shopping Malls," Journal of the Academy of Marketing Science, 15 (Spring): 46-53.
11. Kalwani, Manohar and Alvin J. Silk (1980), “Structure of Repeat Buying for New Packaged Goods,” Journal of Marketing Research, 17 (August), 316-322.
12. Morrison, Donald G. and David C. Schmittlein (1988), “Generalizing the NBD Model for Customer Purchases: What Are the Implications and Is It Worth the Effort?” Journal of Business & Economic Statistics, 6 (2), 1.
13. Reinartz, Werner J. and V. Kumar (2000), “On the Profitability of Long-Life Customers in a Noncontractual Setting: An Empirical Investigation and Implications for Marketing,” Journal of Marketing, 64 (4), 17.
14. Schmittlein, David C., Donald G. Morrison, and Richard Colombo (1987), “Counting Your Customers: Who Are They and What Will They Do Next?” Management Science, 33, 1-24.

APPENDIX.

Moment-Matching Approximation of the Product of Two Gamma Distributions

If x and y are two gamma distributed random variables,

$x \sim$ Gamma (r, a)

$y \sim$ Gamma (s, b)

then the product, $z = xy$, can be assumed to be a gamma distributed random variable

$z \sim$ Gamma (R, A)

with shape and scale parameters, R and A, such that the first two raw moments of the z-distribution is the product of the moments of the x- and y-distributions.

$$m_1^x = \frac{r}{\alpha} \quad m_2^y = \frac{r(r+1)}{\alpha^2}$$

$$m_1^y = \frac{s}{\beta} \quad m_2^y = \frac{s(s+1)}{\beta^2}$$

$$m_2^z = m_1^x . m_1^y = \frac{rs}{\alpha\beta} \quad m_2^z = m_2^x . m_2^y = \frac{r(r+1)s(s+1)}{\alpha^2\beta^2}$$

Since the first moment of the z-distribution, m_1^z , is R/A and the second moment, m_2^z z, is R (R+1)/A^2, we can solve for R and A with the following two equations:

$$\frac{R}{A} = \frac{rs}{\alpha\beta} \quad \frac{R(R+1)}{A^2} = \frac{r(r+1)s(s+1)}{\alpha^2\beta^2}$$

Therefore, the gamma distribution describing the product of two independently distributed gamma random variables has shape and scale parameters that can be calculated from the parameters of the multiplying distributions.

$$R = \frac{rs}{(r+1)(s+1) - rs} \quad A = \frac{\alpha\beta}{(r+1)(s+1) - rs}$$

with Bayesian updating after observing one arrival at time t…

$$R = \frac{(r+1)s}{(r+2)(s+1) - (r+1)s} \qquad A = \frac{(\alpha + t)\beta}{(r+2)(s+1) - (r+1)s}$$

6/22/2018