

# A Simple Probability Model for Projecting Customer Retention

Peter S. Fader  
Bruce G. S. Hardie<sup>1</sup>

September 2005

<sup>1</sup>Peter S. Fader is the Frances and Pei-Yuan Chia Professor of Marketing at the Wharton School of the University of Pennsylvania (address: 749 Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104-6340; phone: 215.898.1132; email: [faderp@wharton.upenn.edu](mailto:faderp@wharton.upenn.edu); web: [www.petefader.com](http://www.petefader.com)). Bruce G. S. Hardie is Associate Professor of Marketing, London Business School (email: [bhardie@london.edu](mailto:bhardie@london.edu); web: [www.brucehardie.com](http://www.brucehardie.com)). The authors thank Michael Berry and Gordon Linoff for providing the data used in this paper, and Naufel Vilcassim for his helpful comments. The second author acknowledges the support of the London Business School Centre for Marketing and the hospitality of the Department of Marketing at the University of Auckland Business School.

## Abstract

### A Simple Probability Model for Projecting Customer Retention

At the heart of any contractual or subscription-oriented business model is the notion of the retention rate. An important managerial task is to take a series of past retention numbers for a given group of customers and project them into the future in order to make more accurate predictions about customer tenure, lifetime value, and so on. In this paper we reanalyze data from a leading book on data mining (Berry and Linoff 2004), who drew the dire conclusion that “parametric approaches do not work” for such a task. As an alternative to common “curve-fitting” regression models, we develop and demonstrate a probability model with a well-grounded “story” for the churn process. We show that our basic model (known as a “shifted-beta-geometric”) can be implemented in a simple Microsoft Excel spreadsheet and provides remarkably accurate forecasts and other useful diagnostics about customer retention. We provide a detailed appendix covering the implementation details and offer additional pointers to other related models.

**Keywords:** retention, churn, forecasting, customer base analysis, probability models, beta-geometric

# 1 Introduction

A defining characteristic of a contractual or subscription business setting is that the departure of a customer is observed. For example, the customer has to contact the firm to cancel her mobile phone contract; similarly, a local theater company can observe that a patron has not renewed his annual subscription.<sup>1</sup> As such, it makes sense to talk of metrics such as retention and churn rates: the retention rate for period  $t$  ( $r_t$ ) is defined as the proportion of customers active at the end of period  $t - 1$  who are still active at the end of period  $t$ , while the churn rate for a given period is defined as the proportion of customers active at the end of period  $t - 1$  who dropped out in period  $t$ .<sup>2</sup>

As we seek to understand the nature of customer behavior in a contractual setting, it is useful to draw on the survival analysis literature. One particularly useful concept for characterizing the distribution of customer lifetimes is that of the *survivor function*, denoted by  $S(t)$ , which is the probability that a customer has “survived” to time  $t$  (i.e., is still active at  $t$ ). Recalling the definition of a retention rate, it follows that

$$\begin{aligned} S(t) &= r_1 \times r_2 \times \cdots \times r_t \\ &= \prod_{i=1}^t r_i, \end{aligned} \tag{1}$$

which implies

$$r_t = \frac{S(t)}{S(t-1)}. \tag{2}$$

Several quantities of managerial interest can easily be calculated directly from the survivor function. For example, the expected (or average) tenure of a customer is simply the area under the survivor function. In a discrete-time setting, this is computed as

$$\text{expected tenure} = \sum_{t=0}^{\infty} S(t).$$

In light of (1), the standard textbook expression for (expected) customer lifetime value (CLV)

---

<sup>1</sup>This is in contrast to a noncontractual setting, a defining characteristic of which is that the departure of a customer is not observed by the firm. See Section 4 for a discussion of the implications of this characteristic.

<sup>2</sup>Strictly speaking, we should talk of retention and churn probabilities, not rates.

in a contractual setting that (correctly) reflects the phenomenon of nonconstant retention rates,

$$E(CLV) = \sum_{t=0}^{\infty} m \left\{ \prod_{i=1}^t r_i \right\} \left( \frac{1}{1+d} \right)^t,$$

can be written as

$$E(CLV) = \sum_{t=0}^{\infty} m \frac{S(t)}{(1+d)^t}.$$

In a contractual setting, the empirical survivor function  $\widehat{S}(t)$  is simply the proportion of customers acquired at time 0 who are still active at time  $t$ . A major problem in using the empirical survivor function to compute expected tenure or lifetime value is that the observed time horizon is often quite limited. Suppose we observe a particular cohort of customers over their first five years with the firm, which implies we can compute  $\widehat{S}(1), \dots, \widehat{S}(5)$ . (By definition,  $\widehat{S}(0) = 1$ .) The quantity  $\widehat{S}(0) + \dots + \widehat{S}(5)$  is the expected customer lifetime for the members of the cohort over this period. Similarly, we can compute expected CLV during the first five years of a customer's relationship with the firm. However, we would be underestimating the expected tenure and CLV of a new customer as we would be ignoring the remaining life of those customers who are alive at the end of the fifth year. In order to compute the true expected tenure and CLV, we need to be able to project the survivor function beyond the observed time horizon. That is, we need to create estimates of  $S(6), S(7), \dots$  given the data  $\widehat{S}(1), \dots, \widehat{S}(5)$ . This projected survivor function is also needed if we wish to compute the expected residual tenure or lifetime value of an individual who has been a customer for, say, three years.

An obvious approach is to fit some flexible function of time to the observed data. The resulting regression equation can then be used to project the survivor function beyond the range of observations, from which we can compute expected tenure, customer lifetime value, etc. In a popular book on data mining, Berry and Linoff (2004) explore this idea (on pages 392–393); their conclusion regarding the viability of such an exercise is evident in the title of their sidebar discussion: “Parametric approaches do not work”.

The objective of this paper is to present an alternative approach to the problem of projecting the survivor function, one that does “work”. We formulate a probabilistic model of contract duration that is based on a simple story of customer behavior. The resulting model offers useful

diagnostic insights and is very easy to implement using Microsoft Excel.

In the next section, we replicate and extend Berry and Linoff’s analysis. We then present a simple probability model of customer lifetime and demonstrate the value of using a formal model to predict future customer behavior. We conclude with a discussion of several issues that arise from this work.

## 2 Projecting Survival Using Simple Functions of Time

The survival data presented in Table 1 are for two segments of customers (“Regular” and “High End”) for an unspecified subscription-type business. These data are presented in graphical form in Berry and Linoff (2004, Chapter 12). The High End data are used by Berry and Linoff in their examination of parametric approaches to the projection of the survivor function.

Year	% survived	
	Regular	High End
0	100.0%	100.0%
1	63.1%	86.9%
2	46.8%	74.3%
3	38.2%	65.3%
4	32.6%	59.3%
5	28.9%	55.1%
6	26.2%	51.7%
7	24.1%	49.1%
8	22.3%	46.8%
9	20.7%	44.5%
10	19.4%	42.7%
11	18.3%	40.9%
12	17.3%	39.4%

**Table 1:** Observed % customers surviving at least 0–12 years

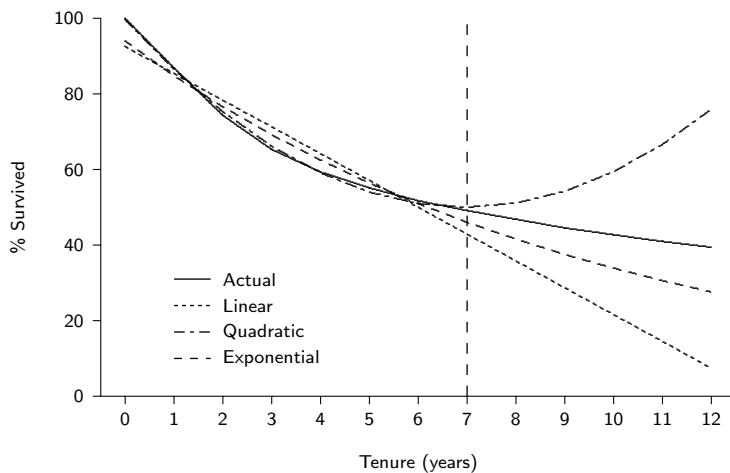
Suppose we only have the first seven years of data and wish to compute estimates of  $S(8), S(9), \dots$ . If we were to give these data to a student who had just completed a typical data analysis course, the natural starting point would be to fit a linear function of time to the data and use the resulting regression equation to project the survivor function out over the future periods. Recognizing that the data are not linear, some students would add a quadratic term to try to capture the curvature in the data. More sophisticated students would specify

some nonlinear function of time, such as an exponential function.

In their “Parametric approaches do not work” sidebar, Berry and Linoff estimate and compare this set of regression models with the following results:<sup>3</sup>

Linear	$y = 0.925 - 0.071t$	$R^2 = 0.922$
Quadratic	$y = 0.997 - 0.142t + 0.010t^2$	$R^2 = 0.998$
Exponential	$\ln(y) = -0.062 - 0.102t$	$R^2 = 0.963$

where  $y$  is the proportion of customers surviving at least  $t$  years. These equations are then used to extrapolate the survivor function out to year 12; Figure 1 re-creates the plot presented in Berry and Linoff’s sidebar (p. 393).



**Figure 1:** Actual versus model-based estimates of the percentage of High End customers surviving at least 0–12 years

The fit of all three models up to and including year 7 is reasonable, and the quadratic model provides a particularly good fit. But when we consider the projections beyond the model calibration period, all three models break down dramatically. The linear and exponential models underestimate year 12 survival by 81% and 30%, respectively, while the quadratic model overestimates year 12 survival by 92%. Furthermore, the models lack logical consistency: the linear model would have  $S(t) < 0$  after year 14, and according to the quadratic model the survival will


<sup>3</sup>In the models run by Berry and Linoff, time is indexed  $1, 2, \dots, 8$ , but in order to maintain consistency with the definitions of  $S(t)$  discussed earlier (specifically  $S(t) = 0$ ), we reindex time to  $0, 1, \dots, 7$ . This has no impact at all on the fit or forecasting performance of any of the models.

start to increase over time, which is not possible. It is therefore not surprising that Berry and Linoff conclude that parametric curves do not “work” for the task of projecting the survivor function over time.

Repeating this analysis for the Regular segment yields the following equations:

Linear	$y = 0.773 - 0.092t$	$R^2 = 0.776$
Quadratic	$y = 0.930 - 0.249t + 0.022t^2$	$R^2 = 0.960$
Exponential	$\ln(y) = -0.248 - 0.190t$	$R^2 = 0.915$

and the corresponding fits and projections are reported in Figure 2. The projections associated with the linear and quadratic models are terrible and illogical once again. The exponential model doesn’t appear to be very bad in the figure, but in fact it underestimates year 12 survival by 54%. This is not an acceptable range of error.



**Figure 2:** Actual versus model-based estimates of the percentage of Regular customers surviving at least 0–12 years

Of course we could try out different arbitrary functions of time but this would be a pure curve-fitting exercise at its worst. Furthermore, it is hard to imagine that there would any underlying rationale for the equation(s) that we might settle upon. Faced with this situation, it is tempting to throw up our hands in despair and say that we cannot project out the survivor function beyond the range of observations.

However, we feel that such a conclusion is premature. After all, in other areas of marketing there are plenty of models that have been used to provide accurate forecasts of the behavior of

a cohort of customers beyond the range of observations (see, for instance, Hardie, Fader and Wisniewski (1998) for the case of new product sales forecasting) . With this in mind, the next section sees us formulating a probabilistic model of contract duration that is based on a simple “story” of customer behavior.

### 3 A Discrete-Time Model for Contract Duration

Consider the following story of customer behavior in a contractual setting:

- i. At the end of each period, a customer flips a coin: “heads” she cancels her contract, “tails” she renews it.
- ii. For a given individual, the probability of a coin coming ups “heads” does not change over time.
- iii.  $P(\text{“heads”})$  varies across customers.

Of course people do not make their contract renewal decisions on the basis of coin flips; rather, this story is a paramorphic representation of customer behavior. The third element of the story should be not be controversial, as the notion of heterogeneity is central to marketing. However, some readers might find the second element contrary to their expectation that retention rates increase over time as the customer gains more experience with the product or service. But rather than overcomplicate our story, we start with the simplest possible set of assumptions and only add supposed richer “touches of reality” if the model does not “work”. As we will see shortly, no additional assumptions will be required in this particular case.

To operationalize this verbal model, we need to translate the elements of this story into the language of mathematics. More formally, our proposed model for the duration of customer lifetimes is based on the following two assumptions:

- i. An individual remains a customer of the firm with constant retention probability  $1 - \theta$ . This implies that the duration of the customer’s relationship with the firm, denoted by the random variable  $T$ , is characterized by the (shifted) geometric distribution with probability

mass function and survivor function

$$P(T = t | \theta) = \theta(1 - \theta)^{t-1}, \quad t = 1, 2, 3, \dots \quad (3)$$

$$S(t | \theta) = (1 - \theta)^t, \quad t = 1, 2, 3, \dots \quad (4)$$

ii. Heterogeneity in  $\theta$  follows a beta distribution with pdf

$$f(\theta | \alpha, \beta) = \frac{\theta^{\alpha-1}(1 - \theta)^{\beta-1}}{B(\alpha, \beta)},$$


where  $B(\cdot, \cdot)$  is the beta function.

The assumption of geometrically-distributed lifetimes follows from the first two elements of our simple story of customer behavior; it is perfectly consistent with the sequential coin-flip description. The beta distribution is less familiar to most readers, but it is a very reasonable way to characterize heterogeneity in the churn probabilities because it is a flexible distribution which is bounded between zero and one. If one thinks about how the “coin-flip” probabilities are likely to vary across individuals, there are four principal possibilities, as illustrated in Figure 3. If both parameters of the beta distribution ( $\alpha$  and  $\beta$ ) are small (less than 1), then the mix of churn probabilities is “U-shaped,” or highly polarized across customers. If both parameters are relatively large ( $\alpha, \beta > 1$ ), then the probabilities are fairly homogeneous. Likewise, the distribution of probabilities can be “J-shaped” or “reverse-J-shaped” if the parameters fall within the remaining ranges as shown in the figure. It is not essential for the reader to remember all of these cases, but these parameters can offer useful diagnostics to help the manager understand the degree (and nature) of heterogeneity in churn probabilities across the customer base.

Since  $\theta$  is unobserved, we take the expectation of (3) and (4) over the distribution of  $\theta$  to arrive at the corresponding expressions for a randomly-chosen individual:

$$P(T = t | \alpha, \beta) = \frac{B(\alpha + 1, \beta + t - 1)}{B(\alpha, \beta)}, \quad t = 1, 2, \dots \quad (5)$$

$$S(t | \alpha, \beta) = \frac{B(\alpha, \beta + t)}{B(\alpha, \beta)}, \quad t = 1, 2, \dots \quad (6)$$



**Figure 3:** General shapes of the beta distribution as a function of  $\alpha$  and  $\beta$

(See Appendix A for details of the derivations.) We call this model the shifted-beta-geometric (sBG) distribution. It has been used in many different settings ranging from familiar business contexts (e.g., consumer response to direct mail solicitations—Buchanan and Morrison (1988)) to very unconventional modeling situations (e.g., the number of menstrual cycles required to achieve pregnancy—Weinberg and Gladen (1986)).

Some readers may be daunted by the presence of beta functions in the above expressions. However it turns out that we can use this model without ever having to deal with a beta function. In Appendix A we show that we can compute sBG probabilities by using the following forward-recursion formula from  $P(T = 1)$ :

$$P(T = t) = \begin{cases} \frac{\alpha}{\alpha + \beta} & t = 1 \\ \frac{\beta + t - 2}{\alpha + \beta + t - 1} P(T = t - 1) & t = 2, 3, \dots \end{cases} \quad (7)$$

Recall from (2) that the retention rate is the ratio of sequential values of the survivor function. Substituting (6) into (2) and simplifying gives us the following expression for the (aggregate) retention rate associated with sBG model:


$$r_t = \frac{\beta + t - 1}{\alpha + \beta + t - 1}. \quad (8)$$

(See Appendix A for details of the derivation.) Given (8), we can compute  $S(t)$  without having to deal with a beta function by using the expression given in (1).

We immediately see that, under the sBG model, the retention rate is an increasing function of time, even though the underlying (unobserved) individual-level retention probability is constant. According to this model, there are no underlying time dynamics at the level of the individual customer; the observed phenomenon of retention rates increasing over time is simply due to heterogeneity (i.e., the high churn customers drop out early in the observation period, with the remaining customers having lower churn probabilities). This well-known “ruse of heterogeneity” (Vaupel and Yashin 1985) is often overlooked by those attempting to make sense of various aggregate patterns of customer behavior.


We fit the sBG model to the first seven years of the data presented in Table 1. For the High End segment,  $\hat{\alpha} = 0.688, \hat{\beta} = 3.806$ ; for the Regular segment,  $\hat{\alpha} = 0.704, \hat{\beta} = 1.182$ . (See Appendix B for details of how to estimate the model parameters in the familiar Microsoft Excel environment.) Using these parameter estimates, we extrapolate the survivor function for each segment out to year 12. These model-based numbers are plotted in Figure 4, along with the corresponding empirical survivor functions. The resulting predictions are almost too good to be true; the sBG model overestimates year 12 survival by only 4% and 2% for the High End and Regular segments, respectively. Even though this model is no more complicated than the regression models discussed earlier, its carefully constructed “story” makes it possible to tease out, and therefore accurately project, the critical behavioral components.

Another plot of interest shows the (aggregate) retention rate as a function of tenure. The model-based retention rate numbers (as computed using (8)) are plotted in Figure 5, along with the corresponding observed retention rates as computed from the empirical survivor functions.



**Figure 4:** Actual versus model-based estimates of the percentage of customers surviving at least 0–12 years for the High End and Regular segments.

For both segments, the sBG model accurately tracks the empirical retention rate curves. On one hand, this might not seem surprising since  $r_t$  and  $S(t)$  are so closely related; on the other hand, however,  $r_t$  is harder to predict accurately since it is not a cumulative number like  $S(t)$  and therefore it is more sensitive to period-to-period variations. Despite the existence of certain unexplained “blips” as in year 2 for the High End segment, the tracking/prediction plot for  $r_t$  is very impressive through year 12 and there is every reason to believe that the model would continue to perform well over an even longer future horizon.




**Figure 5:** Actual versus model-based estimates of retention rates by tenure for the High End and Regular segments.

For both segments we note that the retention rates are an increasing function of the length of a customer’s relationship with the firm. The important point to emphasize, once again, is that

the sBG “story” assumes that these apparent dynamics are simply a result of heterogeneity; any given individual has a constant (but unknown) retention probability  $1 - \theta$ . Unlike the conventional wisdom about customer retention, it is *not* a story of individual customers becoming increasingly loyal as they develop a deeper relationship with the firm, etc.

As a final demonstration of the usefulness of the sBG model, we show and contrast the mixing distributions that characterize how the churn probabilities ( $\theta$ ) differ across the individuals in each segment. In Figure 6 we see that both distributions are “reverse J-shaped.” This implies that, within each group, most customers have fairly low churn probabilities, but there is a sizeable sub-segment within each one that will tend to depart very quickly. These patterns suggest that there is a fairly high degree of heterogeneity within each segment, and therefore a model that doesn’t take these cross-customer differences into account will not perform very well, particularly in terms of out-of-sample forecasting. Closer examination shows that the overall “weight” of the distribution for the Regular group is shifted slightly to the right compared to the High End distribution. This reflects the fact that the Regular group has a higher mean churn probability ( $E(\theta) = \alpha/(\alpha + \beta) = 0.37$ ) compared to that of the High End group ( $E(\theta) = 0.15$ ). It should be clear from Figures 4 and 5 that this kind of difference in the means exists, but this plot provides a better idea about the nature of these differences at a more fine-grained level.



**Figure 6:** Estimated distributions of churn probabilities for the High End and Regular segments

## 4 Discussion

We have presented the shifted-beta-geometric (sBG) distribution as a model for the duration of customer relationships in a contractual setting. This easy-to-implement model enables us to project an empirical survivor function beyond the observed time horizon. As a result, we need not suffer the “truncation” problem associated with computing expected tenure or customer lifetime value using just the observed survival data; that is, underestimating the quantity of interest by ignoring the remaining lifetime of those customers who are still active at the end of the observation period.

Strictly speaking, the sBG model is only applicable in discrete-time contractual settings. In situations where a customer can become inactive at any point in time (rather than only at discrete contract renewal times), it may be more appropriate to use the model’s continuous-time analog, the exponential-gamma (EG) distribution (also known as the Lomax distribution or the “Pareto distribution of the second kind”). This model assumes that the duration of an individual customer’s relationship with the firm is characterized by the exponential distribution, and that heterogeneity in “departure rates” is captured by a gamma distribution (Hardie et al. 1998; Morrison and Schmittlein 1980).

Both the sBG and EG models are based on the assumption that the commonly observed phenomenon of increasing retention rates is due entirely to heterogeneity; individual-customer-level retention rates are assumed to be constant. If we wish to allow for the possibility of time dynamics at the level of the individual customer, we can no longer characterize the duration of an individual’s relationship with the firm using either the shifted-geometric or exponential distribution, both of which have the “memoryless” property (i.e., the probability of survival to  $s + t$ , given survival to  $t$ , is the same as the initial probability of survival to  $s$ ).

A natural extension would be to assume that individual lifetimes can be characterized by the Weibull distribution, which allows for an individual’s risk of cancelling his contract to increase or decrease as the length of the relationship with the firm increases. In a discrete-time contractual setting, this leads to the beta-discrete-Weibull (BdW) model (Fader and Hardie 2005), which is a generalization of the sBG model, while in a continuous-time contractual setting, this leads to a

generalization of the EG model, the Weibull-gamma (WG) model (Hardie et al. 1998; Morrison and Schmittlein 1980; Schweidel et al. 2005).

The latter paper cited above shows how to add additional “bells and whistles” to the retention model, including marketing mix effects, time-varying covariates (such as seasonality), and cross-cohort differences. The key is to bring in all of these factors at the right level, i.e., at the level of the latent parameter of interest (in this case  $\theta$ ) instead of just “jamming” different covariate effects into a regression-like model. Furthermore, we strongly recommend adherence to Occam’s Razor (often expressed as “entities should not be multiplied unnecessarily”), an implication of which is that one should not make any more assumptions than the minimum needed. An appeal to simplicity is particularly important when the managerial focus is more on projection than detailed explanation.

We recognize that the model presented in this paper (along with the extensions discussed above) only apply in a contractual setting, a defining characteristic of which is that the departure of a customer is observed. This is in contrast to a noncontractual setting, a defining characteristic of which is that the departure (or “death”) of a customer is not observed by the firm. In a noncontractual setting, the time at which a customer becomes inactive, and the likelihood that it has occurred at all, must be inferred from the transaction history. This creates many challenges that make the model-building process a tougher task than for the present (contractual) case. Such models are developed by Fader et al. (2005a, 2005b) and Schmittlein et al. (1987). But in those situations, the notion of customer retention is not a meaningful concept anyway, since the company never “owns” the customer, or has any kind of formal relationship that would require renewal. It is important therefore, for managers to use the word “retention” more carefully than current practice often shows. But when retention is indeed a relevant notion, as it is for the dataset used here, then the proposed model and surrounding discussion should be among the first analytical considerations undertaken by management in their desire to better understand and project the retention patterns they have observed.

## Appendix A: Steps in Model Derivation

In this appendix we walk through the derivation of the key mathematical results presented in this paper. It is assumed that the reader is comfortable with the basics of integration.

Both the sBG probability mass function (pmf) and survivor function were originally expressed in terms of beta functions. The beta function  $B(\alpha, \beta)$  is defined by the integral

$$B(\alpha, \beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1} dt, \quad \alpha > 0, \beta > 0. \quad (\text{A1})$$

The beta function can be expressed in terms of gamma functions:

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}.$$

For the purposes of this paper, the only thing we need to know about the gamma function is its so-called recursive property:  $\Gamma(x) = (x-1)\Gamma(x-1)$ .

We derive the expression for the sBG probability mass function in the following manner. If  $\theta$  were known, the probability of dropping out in period  $t$  would simply be the geometric probability  $\theta(1-\theta)^{t-1}$ . But since  $\theta$  is unobserved,  $P(T=t)$  for a randomly-chosen individual is the expected value of the shifted-geometric probability of dropping out in period  $t$  (conditional on  $\Theta = \theta$ ), where the expectation is with respect to the beta distribution for  $\Theta$ :

$$\begin{aligned} P(T=t) &= \int_0^1 P(T=t | \Theta = \theta) f(\theta) d\theta \\ &= \int_0^1 \theta(1-\theta)^{t-1} \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha, \beta)} d\theta \end{aligned}$$

which, combining terms and moving all non- $\theta$  elements to the left of the integral sign,

$$= \frac{1}{B(\alpha, \beta)} \int_0^1 \theta^\alpha (1-\theta)^{\beta+t-2} d\theta.$$

Looking closely at the integral, we see that it is simply the integral expression for the beta

function (A1) with parameters  $\alpha + 1$  and  $\beta + t - 1$ . Therefore,

$$P(T = t) = \frac{B(\alpha + 1, \beta + t - 1)}{B(\alpha, \beta)}.$$

(The expression for the sBG survivor function is derived in exactly the same manner.)

The forward-recursion formula used to compute sBG probabilities is derived in the following manner. We first note that

$$P(T = 1 | \alpha, \beta) = \frac{B(\alpha + 1, \beta)}{B(\alpha, \beta)}$$

which, expressing the beta functions in term of gamma functions and cancelling terms,

$$= \frac{\Gamma(\alpha + 1)}{\Gamma(\alpha)} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha + \beta + 1)}$$

which, recalling the recursive nature of the gamma function,

$$= \frac{\alpha}{\alpha + \beta}.$$

But how does this help us compute  $P(T = t)$  for  $t = 2, 3, \dots$ ? Reflecting on the identity

$$P(T = t) = \frac{P(T = t)}{P(T = t - 1)} \times P(T = t - 1),$$

if we have a simple expression for the ratio  $P(T = t)/P(T = t - 1)$ , we can easily compute  $P(T = 2)$  given the value of  $P(T = 1) = \alpha/(\alpha + \beta)$ . Given the value of  $P(T = 2)$ , we can then compute  $P(T = 3)$ , and so on.

Recalling the sBG pmf, we have

$$\begin{aligned} \frac{P(T = t)}{P(T = t - 1)} &= \frac{B(\alpha + 1, \beta + t - 1)}{B(\alpha, \beta)} \bigg/ \frac{B(\alpha + 1, \beta + t - 2)}{B(\alpha, \beta)} \\ &= \frac{B(\alpha + 1, \beta + t - 1)}{B(\alpha + 1, \beta + t - 2)} \end{aligned}$$

which, expressing the beta functions in term of gamma functions and cancelling terms,

$$= \frac{\Gamma(\beta + t - 1) \Gamma(\alpha + \beta + t - 1)}{\Gamma(\beta + t - 2) \Gamma(\alpha + \beta + t)}$$

which, recalling the recursive nature of the gamma function,

$$= \frac{\beta + t - 2}{\alpha + \beta + t - 1}.$$

The complete forward-recursion formula naturally follows.

Finally, to derive the expression for the retention rate as implied by the sBG model, we substitute the expression for the sBG survivor function into (2) and simplify:

$$\begin{aligned} r_t &= \frac{B(\alpha, \beta + t)}{B(\alpha, \beta)} \bigg/ \frac{B(\alpha, \beta + t - 1)}{B(\alpha, \beta)} \\ &= \frac{B(\alpha, \beta + t)}{B(\alpha, \beta + t - 1)} \end{aligned}$$

which, expressing the beta functions in term of gamma functions and cancelling terms,

$$= \frac{\Gamma(\beta + t) \Gamma(\alpha + \beta + t - 1)}{\Gamma(\beta + t - 1) \Gamma(\alpha + \beta + t)}$$

which, recalling the recursive nature of the gamma function,

$$= \frac{\beta + t - 1}{\alpha + \beta + t - 1}.$$

## Appendix B: Implementing the Model in Excel

In this appendix we show how to compute the maximum likelihood estimates the sBG model parameters for the High End dataset using Microsoft Excel. Before providing step-by-step instructions for constructing the worksheet, we briefly review the notion of maximum likelihood estimation.

Suppose we observe a group of  $n$  customers for seven periods. We note that  $n_1$  customers drop out in the first period,  $n_2$  in the second period,  $\dots$ , with  $n_7$  customers departing in the seventh period. It follows that  $n - \sum_{t=1}^7 n_t$  customers are still active at the end of the seventh period.

Let us assume that the customer lifetimes can be characterized by the sBG distribution. What is the probability that a randomly-chosen customer has a lifetime of one period? The answer is the sBG probability  $P(T = 1 | \alpha, \beta)$ . What is the probability that a randomly-chosen customer has a lifetime of two periods? Answer: the sBG probability  $P(T = 2 | \alpha, \beta)$ . What is the probability that one randomly-chosen customer has a lifetime of one period while another has a lifetime of two periods? Assuming the propensity of one customer to drop out is independent of the behavior of the other customer, it is simply the product of the respective sBG probabilities:  $P(T = 1 | \alpha, \beta)P(T = 2 | \alpha, \beta)$ . It follows that, given specific values of the model parameters  $\alpha$  and  $\beta$ , the joint probability of  $n_1$  customers departing in the first period,  $n_2$  in the second period,  $\dots$ ,  $n_7$  in the seventh period, and  $n - \sum_{t=1}^7 n_t$  customers still being active at the end of the seventh period is

$$\begin{aligned}
 P(\text{data} | \alpha, \beta) &= P(T = 1 | \alpha, \beta)^{n_1} P(T = 2 | \alpha, \beta)^{n_2} P(T = 3 | \alpha, \beta)^{n_3} \\
 &\quad \times P(T = 4 | \alpha, \beta)^{n_4} P(T = 5 | \alpha, \beta)^{n_5} P(T = 6 | \alpha, \beta)^{n_6} \\
 &\quad \times P(T = 7 | \alpha, \beta)^{n_7} S(7 | \alpha, \beta)^{n - \sum_{t=1}^7 n_t}. \tag{B1}
 \end{aligned}$$

However, we do not know the values of  $\alpha$  and  $\beta$ , even though we believe that the data come from the sBG distribution.

The idea of maximum likelihood estimation is to ask what values of the model parameters

maximize the probability (or, more formally, the *likelihood*) of the observed data. We define the likelihood function as

$$\begin{aligned}
 L(\alpha, \beta \mid \text{data}) &= P(T = 1 \mid \alpha, \beta)^{n_1} P(T = 2 \mid \alpha, \beta)^{n_2} P(T = 3 \mid \alpha, \beta)^{n_3} \\
 &\quad \times P(T = 4 \mid \alpha, \beta)^{n_4} P(T = 5 \mid \alpha, \beta)^{n_5} P(T = 6 \mid \alpha, \beta)^{n_6} \\
 &\quad \times P(T = 7 \mid \alpha, \beta)^{n_7} S(7 \mid \alpha, \beta)^{n - \sum_{t=1}^7 n_t}. \tag{B2}
 \end{aligned}$$

and use numerical optimization methods (e.g., the “Solver” add-in in Excel) to find the values of  $\alpha$  and  $\beta$  that maximize this function; these are called the *maximum likelihood estimates* of the model parameters.<sup>4</sup> As the number computed using (B2) will be very small, we usually work with the natural logarithm of the likelihood function, the so-called log-likelihood function:

$$LL(\alpha, \beta \mid \text{data}) = \sum_{t=1}^7 n_t \ln [P(T = t \mid \alpha, \beta)] + \left( n - \sum_{t=1}^7 n_t \right) \ln [S(7 \mid \alpha, \beta)]. \tag{B3}$$

The observant reader will note that we do not actually know  $n, n_1, n_2, \dots, n_7$  for the two datasets given in Table 1; the data are expressed as percentages of the initial number of customers. Looking closely at (B3), we see that this is not a problem; we can simply factor out  $n$  (e.g.,  $n_1$  becomes  $n_1/n$ , the proportion of customers who become inactive in the first period). While this will affect the height of the function, the location of the maximum (i.e., the values of  $\alpha$  and  $\beta$ ) will be unaffected.

So our task is to “code up” this expression for the model log-likelihood function in an Excel worksheet and find maximum likelihood estimates of  $\alpha$  and  $\beta$  by using Solver to find the values of  $\alpha$  and  $\beta$  that maximize the value of this function. The relevant worksheet is shown in Figure B1 and is constructed in the following manner.

- In order to enter expressions for  $P(T = t \mid \alpha, \beta)$  without an error message appearing (e.g., #NUM! or #DIV/0!), we need some “starting values” for  $\alpha$  and  $\beta$ . The exact values do not matter—provided they are within the defined bounds—so we start with 1.0 for  $\alpha$  and  $\beta$ , locating these parameter values in cells B1:B2, respectively.

---

<sup>4</sup>We note that (B1) and (B2) look almost identical, but there is a subtle difference: in (B1), the probability we compute is a function of the data pattern for fixed model parameters, while in (B2), we already have the data and the probability we compute is a function of the model parameters.

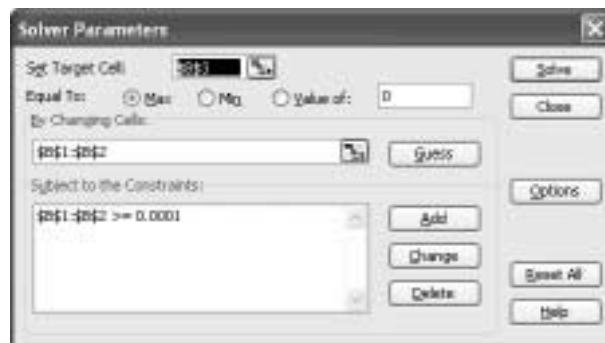
	A	B	C	D	E	F
1	alpha	1.000				
2	beta	1.000				
3	LL	-2.116				
4						
5	t	P(T=t)	S(t)	% alive	% die	
6	1	0.500	0.500	86.9%	13.1%	-0.091
7	2	0.167	0.333	74.3%	12.6%	-0.226
8	3	0.083	0.250	65.3%	9.0%	-0.224
9	4	0.050	0.200	59.3%	6.0%	-0.180
10	5	0.033	0.167	55.1%	4.2%	-0.143
11	6	0.024	0.143	51.7%	3.4%	-0.127
12	7	0.018	0.125	49.1%	2.6%	-0.105
13						-1.021

**Figure B1:** Screenshot of Excel Worksheet for Parameter Estimation

- We enter the values of  $t = 1, 2, \dots, 7$  in cells A6:A12.
- The corresponding values of  $P(T = t | \alpha, \beta)$  are computed in cells B6:B12 using the forward-recursion given in (7):
  - We compute  $P(T = 1)$  by entering `=B1/(B1+B2)` in cell B6.
  - We compute  $P(T = 2)$  by entering `=(B$2+A7-2)/(B$1+B$2+A7-1)*B6` in cell B7.
  - We copy B7 to B8:B12.
- We compute the values of  $S(t | \alpha, \beta)$  for  $t = 1, 2, \dots, 7$  in cells C6:C12:
  - $S(1)$  is simply  $1 - P(T = 1)$ , so we enter `=1-B6` in cell C6.
  - For  $t > 1$ ,  $S(t) = S(t - 1) - P(T = t)$ , so we enter `=C6-B7` in cell C7.
  - We copy C7 to C8:C12.
- The next step is to enter the actual survival data. The proportion for year 1 (0.869) is entered in cell D6, the proportion for year 2 (0.743) is entered in cell D7, and so on down to 0.491 in cell D12 for year 7. (In the worksheet shown in Figure B1, cells D6:D12 are formatted using the percentage style.)
- The proportion of customers dropping out each year, as required for the log-likelihood function, is computed in cells E6:E12:

- As the proportion of customers who dropped out in year 1 is simply one minus the proportion of customers who are still active at the end of the first year, we enter =1-D6 in cell E6.
- For  $t > 1$ , the proportion of customers who dropped out in year  $t$  is the proportion of customers who are still active at the end of year  $t - 1$  minus the proportion of customers who are still active at the end of the year  $t$ . We therefore enter =D6-D7 in cell E7 and copy it to E8:E12.
- The first seven elements of the log-likelihood function are computed in cells F6:F12: we enter =E6\*LN(B6) in cell F6 and copy it to E7:E12.
- The final element of the log-likelihood function, that associated with those customers who have survived at least seven years, is entered as =D12\*LN(C12) in cell F13.
- The sum of cells F6:F13 is entered in cell B3; this is the value of the log-likelihood function given the values for the two model parameters in cells B1:B2. (With starting values of 1.0 for both parameters,  $LL = -2.116$ .)

We find the maximum likelihood estimates of the two model parameters by maximizing the log-likelihood function. We do this using the Excel add-in Solver, available under the “Tools” menu. The *target cell* is the value of the log-likelihood, cell B3. We wish to *maximize* this by *changing* cells B1:B2. The *constraints* we place on the parameters are that  $\alpha$  and  $\beta$  are greater than 0. As Solver only offers us a “greater than or equal to” constraint, we *add* the constraint that cells B1:B2 are  $\geq$  a small positive number (e.g., 0.0001) — see Figure B2.



**Figure B2:** Solver Settings

Clicking the *Solve* button, Solver converges to a solution where the maximum value of the log-likelihood function is  $-1.611$ , associated with  $\alpha = 0.668$  and  $\beta = 3.806$ . These are the maximum likelihood estimates of the model parameters. (So as to be sure that we have actually reached the maximum of the log-likelihood function, it is good practice to redo the optimization process using a completely different set of starting values. For example, using starting values of  $0.01$  and  $0.01$  (for which  $LL = -2.742$ ), use Solver to find the maximum of the log-likelihood function. Are the corresponding values of the two model parameters equal to those given above? They should be!)

## References

- Berry, Michael J. A. and Gordon S. Linoff (2004), *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management*, 2nd edition, Indianapolis, IN: Wiley Publishing, Inc.
- Buchanan, Bruce and Donald G. Morrison (1988), "A Stochastic Model of List Falloff with Implications for Repeat Mailings," *Journal of Direct Marketing*, **2** (Summer), 7–15.
- Fader, Peter S. and Bruce G. S. Hardie (2005), "Accommodating Individual-level Dynamics in a Discrete Lifetime Distribution," unpublished working paper.
- Fader, Peter S., Bruce G. S. Hardie, and Ka Lok Lee (2005a), "Counting Your Customers" the Easy Way: An Alternative to the Pareto/NBD Model," *Marketing Science*, **24** (Spring), 275–284.
- Fader, Peter S., Bruce G. S. Hardie, and Ka Lok Lee (2005b), "RFM and CLV: Using Iso-value Curves for Customer Base Analysis," *Journal of Marketing Research*, **42** (November).
- Hardie, Bruce G. S., Peter S. Fader, and Michael Wisniewski (1998), "An Empirical Comparison of New Product Trial Forecasting Models," *Journal of Forecasting*, **17** (June–July), 209–229.
- Morrison, Donald G. and David C. Schmittlein (1980), "Jobs, Strikes, and Wars: Probability Models for Duration," *Organizational Behavior and Human Performance*, **25** (April), 224–251.
- Schmittlein, David C., Donald G. Morrison, and Richard Colombo (1987), "Counting Your Customers: Who They Are and What Will They Do Next?" *Management Science*, **33** (January), 1–24.
- Schweidel, David A., Peter S. Fader, Peter and Eric T. Bradlow (2005), "Modeling Retention in and Across Cohorts," <http://ssrn.com/abstract=742884>.
- Vaupel, James W. and Anatoli I. Yashin (1985), "Heterogeneity's Ruses: Some Surprising Effects of Selection on Population Dynamics," *The American Statistician*, **39** (August), 176–185.
- Weinberg, Clarice Ring and Beth C. Gladen (1986), "The Beta-Geometric Distribution Applied to Comparative Fecundability Studies," *Biometrics*, **42** (September), 547–560.