Ehrenberg-Bass Institute **Working Paper**:

New approaches to estimating NBD-Dirichlet model parameters, from measuring goodness-of-fit using Euclidean Distance

This working paper, dated 27 December 2023, is forthcoming in the ***Social Science Research Network.***

Authors:

Nick Danenberg - Ehrenberg-Bass Institute
Phouthun Nay - Ehrenberg-Bass Institute

University of South Australia | Ehrenberg-Bass Institute for Marketing Science

Technical Note: New approaches to estimating NBD-Dirichlet model parameters, from measuring goodness-of-fit using Euclidean Distance.

Nick Danenberg and Phouthun Nay

Abstract

The NBD-Dirichlet Model is used in marketing to model buying behaviour in competitive, repeat-purchase markets and to estimate brand performance measures. For the marketing practitioner, comparison of the actual, observed in-market values with theoretical values produced by the model suggests the model generalises across a very wide range of category and brand contexts. However, even the developers of the original model have long called for ongoing endeavours to continually improve upon the measure of model fit. In this paper, we put forward an alternative approach for estimation of the parameters to the model, which produces a model that corrects for a systematic error in the estimates for large brands and improves fit of the model overall.

Background

The NBD-Dirichlet Model (also referred to by the moniker the "N-D Model", Goodhardt, Ehrenberg and Chatfield, 1984) was developed for use in a marketing context several decades ago by Goodhardt, Ehrenberg and Chatfield (Ehrenberg 1959; Goodhardt, Ehrenberg and Chatfield, 1984). The NBD-Dirichlet mathematically describes buying behaviour in competitive, repeat-purchase markets, by modelling the patterns of repeat purchases of the brands among consumers within a product category.

Since its development, the NBD-Dirichlet has been applied to a vast number of markets with considerable success in terms of describing what the structure of consumers' repeat purchasing from among a repertoire of brands will look like (Ehrenberg, Uncles and Goodhardt, 2004). The widespread success of the NBD-Dirichlet over decades of application in a vast number of categories has led to the model becoming regarded as an empirical generalisation, one of only a few in marketing (Bass, 1995; Rungie and Goodhardt, 2004; Bound, 2009; Uncles, Ehrenberg and Hammond, 1995; Ehrenberg, Uncles and Goodhardt, 2004).

NBD-Dirichlet modelling

The NBD-Dirichlet is particularly useful to marketing academics and practitioners alike with regards to the modelling and analysis of what are known as 'brand performance measures' (Ehrenberg, 1988; Uncles, Ehrenberg and Hammond, 1995; Ehrenberg, Uncles and Goodhardt, 2004), such as, market share, average purchase frequency, penetration, share of category requirements, 100% loyals, average portfolio size and repeat rate. When undertaken for the purpose of generating estimates of such brand performance measures for each brand, theoretical estimates are compared to the respective observed values derived directly from the data (Ehrenberg, Uncles and Goodhardt, 2004). The high performance of the NBD-Dirichlet, in terms of the accuracy of fit of the model's theoretical output values to observed 'real-world' data, has made the model especially useful as a comparison or benchmark for other modelling approaches (Driesener, Banelis and Rungie, 2017; Kalwani, Meyer and Morrison, 1994).

The process of NBD-Dirichlet modelling begins by taking some observed purchase data within a product category, aggregated from individual-level purchasing (commonly either individual panellists or households) over a fixed period of observation such as a month, a quarter or year, and deriving a fit to theoretical model values from estimating the various model parameters. Fitting the NBD-Dirichlet model is considered as being a fairly complex task to undertake (Ehrenberg, Uncles and Goodhardt, 2004; Rungie and Goodhardt 2004) and remains so even today. This is largely due to the nature of the equations underpinning the model that need to be solved in the estimation process but which have no closed form solutions, relying as they do on the iterative calculations of a number of infinite series. As a result of this complexity, the fitting and the application of the NBD-Dirichlet is still typically carried out through the use of specialised software packages. A number of such packages have been in general circulation over the last quarter-century, released on a variety of platforms: 'Buyer' software (Uncles, 1989), a DOS program that could work from the raw, individual-level data; 'Dirichlet' (Kearns, 2002), an Excel workbook based approach working with the already-aggregated data as inputs, together with a User's Guide written by Bound (2009); 'NBDDirichlet' (Chen, 2008), a package written in and for the R programming language.

Despite the model's comprehensive formal exposition in various publications since its conception in 1984, one of the model's developers wrote some twenty years later that he felt that the steps involved in the estimation procedures still had not then been concisely documented (Rungie and Goodhardt, 2004). In order to support the advancement of the use of the NBD-Dirichlet model, Rungie and Goodhardt (2004) wrote a technical note about fitting the parameters of the NBD-Dirichlet to a set of data inputs. They documented the functional forms, and their derivations, for calculating 'brand performance measures' from fitting the parameters of the NBD-Dirichlet to a set of data inputs within a given category (or, 'product field' as Ehrenberg and colleagues referred to it) and the competing brands that comprise that category. The clear exposition of the calculation of these measures had not hitherto been so concisely documented. Rungie and Goodhardt (2004) were hopeful that in writing their note it would support the ongoing further development and adoption of the NBD-Dirichlet model. A hope that is enacted in the developments that this paper describes.

NBD-Dirichlet Estimation

In broad terms, there are three classes of approach for estimating the parameters of the NBD-Dirichlet Model. The first two are methods that use aggregated data in terms of purchase rates: the method of means and zeros and the method of moments. Both of these methods were espoused and have been used by the developers of the model (Ehrenberg, 1988; Goodhardt, Ehrenberg and Chatfield, 1984) and others since the model's development. The third approach, which was developed later, requires individual-level purchasing data and uses likelihood theory for the estimation modelling. This approach is claimed to offer the benefits of being more efficient, subject to less sampling variation and can even be used to extend the NBD-Dirichlet Model into a generalised model (Rungie and Goodhardt, 2004; Rungie, Laurent and Stern, 2003) through combining with other 'marketing mix' variables. Aggregated purchase data, though, is vastly more readily available (over the raw, individual-level purchase data) to analysts in industry and academia alike. Hence, the former two methods, and of these primarily the method of means and zeros, has been the de-facto standard approach employed in NBD-Dirichlet model fitting. The method of means and zeros is now in widespread use, in particular through the popularity of the widely available 'Dirichlet' Excel workbook (Kearns, 2002). One of the appealing features of the method of means and zeros is that it has the very useful property of generating estimates of the market shares for the brands which exactly match the observed market shares (Rungie and Goodhardt, 2004). This makes it particularly well suited to the practical application of the NBD-Dirichlet Model in the estimation and analysis of brand performance measures e.g., in industry such as testing whether brands within a product market are competing with one another as the model would predict.

Parameter Estimation

In modelling the NBD-Dirichlet using the method of means and zeros, there are three key parameters from which all the other brand performance metrics are derived: the K, A and S parameters. These parameters loosely represent category concentration, weight of purchasing and competition, respectively (Driesener, 2005; Habel and Lockshin, 2013). Of these, only two parameters the K and S parameters, must be uniquely estimated — the parameter A and all other category and brand performance metrics are derived directly from substituting the unique K and S parameter estimates into equations along with other input metrics.

Once the parameter K (and from this, A) has been estimated, the rather more complicated task of deriving an estimate of the S parameter can proceed. The usual approach for this, as laid out in Ehrenberg's treatise "Repeat Buying" (1988) is to iteratively calculate individual S estimates for each brand separately and then combine these into the final category S through calculating the weighted average of individual brand-level S estimates, based on brand size. The specific calculations used to solve for a brand's S are best discussed elsewhere (e.g. Ehrenberg, 1988; Goodhardt, Ehrenberg and Chatfield, 1984; Rungie and Goodhardt, 2004, among others). In practice, when model fitting within typical product categories, there is considerable variation in the brand-level S estimates — the S estimate that provides the best solution for one brand is not necessarily (and in fact, rarely is, under real world conditions) the S estimate that is the best solution for any other brand. Hence, an analytical decision rule must be made about which S estimate to employ as being best representative of the category as a whole, which would then be used to feed through all the subsequent calculations of the category and brand performance metrics. The tractable solution to this problem that the NBD-Dirichlet model developers suggested was to derive an average estimate of S, whereby the S estimate of each brand is taken into consideration, but doing so through applying a weight based on each brand's size (market share). Hence, a (final) <u>category</u> S is constituted from each brand's S estimate, with a greater weighted contribution coming from the larger brands, and a smaller contribution from smaller brands.

Brand-size weighting of S therefore carries with it the risk that the S used in analysis is most highly 'tuned' towards the larger brands, or indeed, the largest brand. NBD-Dirichlet model fitting is an analysis approach that typically will have the aim of identifying outlier, or deviating, brands. For instance, identifying those brands for which their observed brand performance measure values deviate significantly from the model's predicted theoretical values. This is a necessary first step in the extremely useful task of seeking to understand how and why those deviations have come about. In such an approach, using an S metric that is weighted heavily towards the biggest brands, leads to the situation where the deviations that do exist will be more likely to be evident among the smaller brands, as the S that would best fit those brands has contributed only minimally in the calculation of the category S used in analysis. And conversely, the resultant category S is *ipso facto* going to most closely approximate the S that best fits the largest brand/s.

A different method of deriving the category S estimate that does so independently of brand size, treating each brand uniformly, irrespective of size may address the potential for artefact introduced through brand-size weighting, whereby deviations may be more commonly observed among smaller brands. One such approach would be if there is a measure of fit which could then be optimised for; solving S for the value that minimises the error between the observed and the theoretical predicted values. This could be accomplished computationally by iteratively trying a number of values for S, choosing the value that minimises the error.

Current approaches to measuring the fit of the NBD-Dirichlet

Measuring the fit of the NBD-Dirichlet model has been of interest to academic and industry analysts for a long time. See Driesener, Banelis and Rungie (2017) for a comprehensive discussion of this practice. Until now this has universally been undertaken through the investigation of deviations from the observed to expected values for each of the separate brand performance metrics individually (usually penetration and purchase frequency). Driesener, Banelis and Rungie (2017) canvassed the range of measures that have been used to assess the model's fit. They then proposed the use of a suite of uni-dimensional goodness-of-fit statistics for each performance metric. The statistics they proposed are: average percentage error, correlation, mean absolute deviation and mean absolute proportional error. These goodness-of-fit statistics are calculated for each brand performance metric. Each statistic is in turn assessed against benchmark values to derive a suite of pass-fail benchmark criteria. The overall score of fit is the number of tests passed (out of 8) across the suite of tests. This goodness-of-fit suite of tests has subsequently been adopted to become a common practice.

To date, NBD-Dirichlet goodness-of-fit measures have decomposed the fit of the model into separate uni-dimensional measures, such as most commonly the measures of penetration and purchase frequency. The x-y relationship between 'brand size' variables (Penetration or Market Share, for instance) and 'loyalty' metrics (Purchase Frequency, for instance) in the NBD-Dirichlet model can be expressed visually in the form of the well-known"Double Jeopardy (DJ) curve" (or, line). When the DJ relationship is visualised as a curve, it is most commonly shown as a graphical plot of the relationship between Penetration (on the x-axis) and Purchase Frequency (on the y-axis), which demonstrates the curvilinear, upward-sloping relationship between the two variables.

The DJ curve encapsulates the theoretical relationship between two variables (typically penetration and purchase frequency) for a given set of model input parameters that are unique within a given 'product field' of product category and brands. Given the provided inputs, a brand with a given level of penetration will theoretically have a unique purchase frequency of buying, among those customers having bought the brand at all. Hence, a given DJ curve is unique to that 'product field' of analysis with there being a single DJ curve that describes the relationship between penetration and purchase frequency. Another interpretation of the DJ curve is that any individual point along the DJ curve can be considered to be a point of market share (or a basis-point of share, or indeed whatever degree of resolution at which the analyst chooses to calculate the DJ curve), that is unique within that product field. According to the NBD-Dirichlet model, for a given set of input parameters, each point of market share (unit sales), will be uniquely comprised of x% of customers each buying, on average y number of times. Creating a given DJ line is not straightforward, however. This is because due to the nature of the NBD-Dirichlet model, there is no single equation that formally describes the mathematical form of the DJ curve relationship. Habel and Rungie (2013) documented the process of creating a DJ curve from NBD-Dirichlet model parameters and outputs. The process involves iterating through each point of market share (e.g. 100-points) and for each point solving for the corresponding unique penetration and purchase frequency values. Thus, the DJ is a 'perfect' representation of the x-y relationship between the dimensions of penetration and purchase frequency. However the DJ relationship is curvilinear and not a straight line (despite it oftentimes being called a DJ Line). This means that measuring error between the observed and theoretical, predicted values in terms of 'sum of squared errors' is not applicable. Likewise there is not a standard measure for fit, such as an R-squared metric, that could be used to optimise model fit against, i.e. through minimising error.

A single, multi-dimensional goodness-of-fit measure for the NBD-Dirichlet

Until now, therefore the task of deriving a single S parameter estimate that best fits the product field being analysed through optimisation has been intractable. This is because there is currently no single goodness-of-fit measure that would allow optimisation through minimising model error. Consequently, the methods of deriving parameter estimates have remained those that the model developers coined decades ago, viz. iteratively deriving S estimates for individual brands and averaging these to combine into a category-level S estimate.

However, we propose a novel solution to this problem, based on the parsimonious measure of Euclidean Distance that calculates error by accounting for both dimensions jointly. Euclidean Distance is a well-established measure (Habel and Lockshin, 2013) that has long been used in analytical approaches such as correspondence analysis (Greenacre, 1984) and cluster analysis (Aldenderfer and Blashfield, 1984). Euclidean Distance has also been applied to the realm of NBD-Dirichlet modelling (Habel and Lockshin, 2013). Habel and Lockshin (2013) used the Mean Standardised Euclidean Distance (MSED) to measure error in their investigation of which simplified functional form of the Double Jeopardy curve (linear, exponential, simplified curvilinear relationships using the w_0 approximation, and the method of plotting points of market share) yields the best fit to the observed data.

The calculation of ED between an observed point (b_o, w_o) and an estimated point (b_e, w_e), where b and w are Penetration and Purchase Frequency respectively is given as:

$$ED = \sqrt{(w_0 - w_e)^2 + (b_0 - b_e)^2}$$

(Pare, 2000; Habel and Lockshin, 2013).

The axes are first standardised to reduce scale bias in either dimension through dividing each point's values by the respective maximum scale values, to create Standardised Euclidean Distances (SED) for each point. These SED values for each brand in the product category are then averaged to give the Mean Standardised Euclidean Distance (MSED). The MSED is therefore a measure of fit that is directly comparable across product categories and separate analyses, as it is independent of the nature of the category and the number of brands analysed (Habel and Lockshin, 2013). Being an error measure, smaller MSED values indicate a better model fit.

There are two obvious ways that the MSED measure can be calculated. The first, through using the SED values from the two-dimensional observed points (Penetration, Purchase Frequency) to the theoretical points for each respective brand. That is, through essentially calculating the ED between the observed points for each brand of its given market share and the point on the DJ curve where a brand of that given market share would reside, in terms of the Penetration and Purchase Frequency. The second is to merely calculate the ED of each observed point to the nearest point on the DJ curve, irrespective of what market share that point on the DJ curve would equate to. In this way, the "distance to line" approach is more akin to a simple measure of fit to a line, as a person would intuitively 'eyeball' the fit themselves. Using the MSED in this "distance to line" method was previously employed by Habel and Lockshin (2013).

The calculation of the DJ curve itself remains computationally intensive, since it involves the process of iteratively calculating the Penetration and Purchase Frequency (x, y) coordinates for each point of market share. And the process of calculating each point of Penetration involves a number of infinite series, since there is no closed-form solution to calculate penetration from market share. Though then deriving Purchase Frequency once Penetration has been solved is quite rudimentary.

To yield an appropriate level of resolution in the DJ curve, we recommend using at least 1,000 basis-points of market share. Habel and Lockshin worked with 100 points (see their paper for the derivation of a DJ curve). Thus, despite being referred to as a curve, as it certainly is visually, it really exists as a series of discrete points, and not a curve function in a mathematical sense. For these reasons it is not directly suitable to optimise the NBD-Dirichlet fit using this "distance to line" approach. Hence, we advocate using the MSED values using the former approach which calculates the distance from the observed value for each brand to its theoretical value. Thereby calculating a comparatively few number of SEDs and calculating the overall category-level MSED from them. The optimisation routines would then optimise using this MSED metric. The MSED derived from the distance to the DJ curve is still a very useful metric, yielding a measure of the error in the chart relationship. However, we recommend that it be calculated only once, for the model parameters derived through the prior optimisation routines instead of being used iteratively in optimisation.

Comparison of model fits

We now demonstrate the difference in model fit between the various methods of deriving NBD-Dirichlet parameters. We do so using previously published data from a doctoral thesis that investigated a number of product categories in Australia and New Zealand (Sharp, 1997). As a test of fit, we merely display the DJ curve plots to 'eyeball' the difference in fit between the various methods. We don't here look to perform an exhaustive comparison to all product categories, but rather pick just two categories as illustrative examples which are both in the Australian market with the purchasing population of both being Fly Buys membership base only; Fuel and Credit Cards.

Firstly, we compare the fit derived from the brand-size weighted method for estimating S (1) with the model fit derived from optimising the S parameter to best fit all the brands in the product category, for a given K parameter that has first been calculated (2). Then we compare the fits of these methods to a third method, which derives the values for both the K and S parameters jointly, to determine the values for both that yields the best fit to the observed values.

Category 1: Australian Fuel, Fly Buys Members

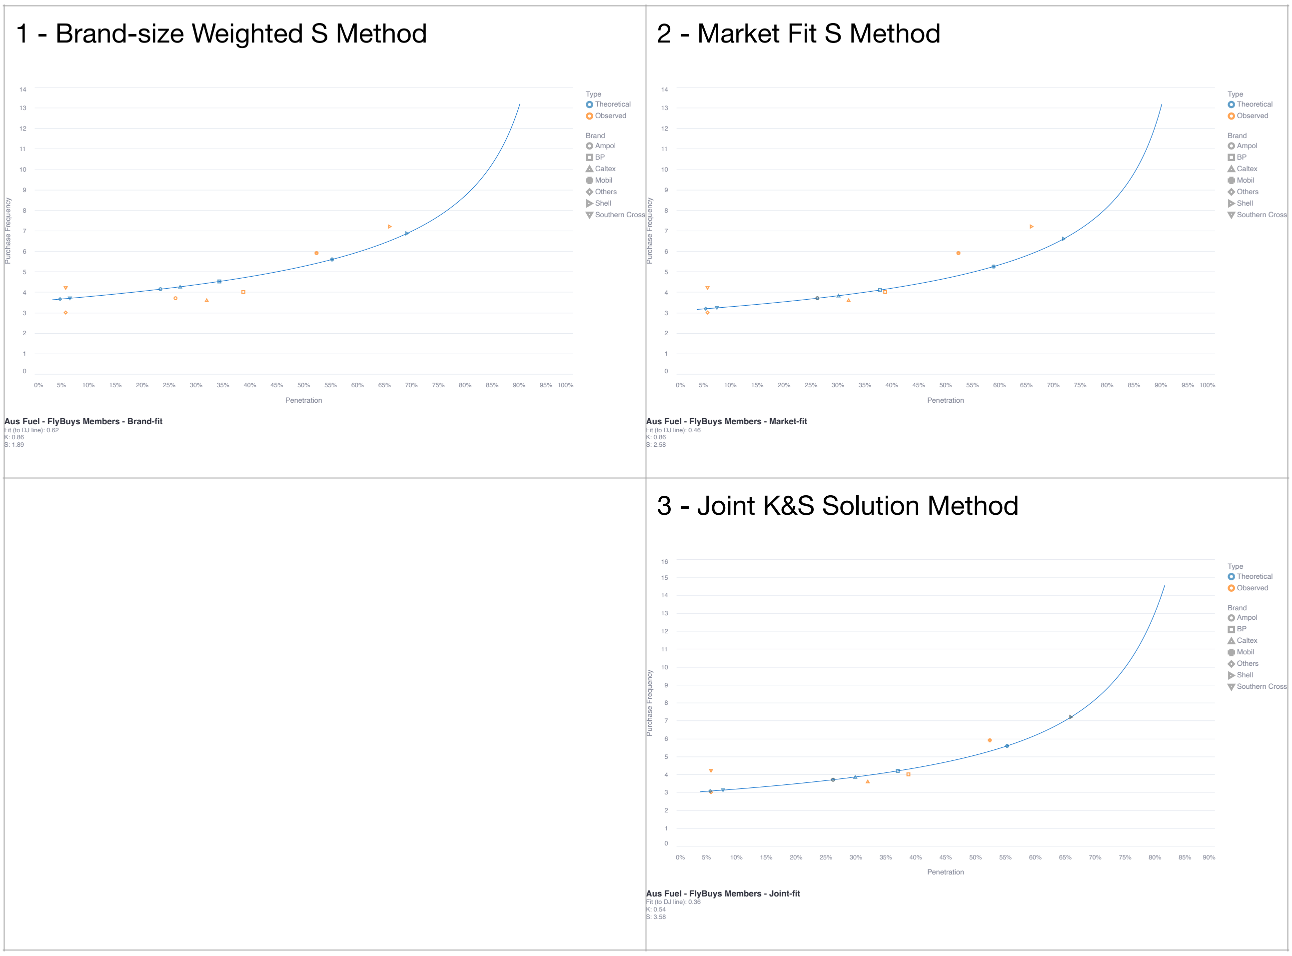

In the brand-size method, we can see that the two largest penetration brands are quite close to the line (and their predicted values), albeit slightly above the line. Although off the line, neither brand would likely be considered to be ‘deviations’. To some extent this has been ‘at the expense’ of the fit for the ‘middle size’ brands (those with penetrations between 25% and 40%) that sit clearly below the line. Although being close to the line in a vertical dimension, the two larger of these brands are some distance away from their theoretical points, and may well then be considered to be deviations in terms of distance between observed and theoretical points. The line passes through the middle of the two very small brands, with both sitting off the line, but probably not far away from their theoretical points to be considered as deviations.

Whereas in the case of the market-fit method, both the two largest brands are further away from the line (in vertical dimension, and also their respective predicted points). And both could be considered as being deviations, both likely being classified as exhibiting 'excess loyalty'. The middle-size brands, on the other hand are now all very close to the line and their respective theoretical points. In fact, the smallest of these mid-size brands has the theoretical point 'perfectly' coinciding with the observed datapoint. The result being that none of the mid-size brands would now be considered as being a 'deviation'. For the smallest brands, the brand that was sitting above the line is now sitting further above it and perhaps might be considered a slight deviation.

Essentially what has happened is that the market-fit method of estimating the S parameter has resulted in the DJ curve being moved down vertically (although the maximum point, the point of 99.99% market share, remains fixed and the curve 'pivots' on this point, though seen as moving downwards for the majority of its length) . And in doing so, has given rise to those effects stated above.

The MSED of the fit the observed data points to the nearest point on the DJ line (in terms of the 1,000 basis points of market share) has decreased from 0.62 to 0.46, indicating a superior overall fit from the market-fit method in comparison to the brand-fit method.

Secondly, we compare the fit of the first two methods with that from the joint-fit method (that jointly derived K and S parameters). The first point to note is that due to a different K parameter being derived and used (0.54 cf 0.86 in the first two methods), this results in a different scaling of the axes. Whereas in the first two methods, in both cases a 99.99% market share brand would be comprised of ≈90% Penetration and 13.2 Average Purchase Frequency (specified from just visual inspection of the charts); with very little difference between the two. Whereas the different K value derived for the joint-fit method results in a 99.99% market share brand would be comprised of ≈82% Penetration and 14.5 Average Purchase Frequency. This rescaling therefore does more than just 'moving' the DJ curve down. It completely rescales the curve. This allows the following relationships between observed and theoretical points to be exhibited. The first thing to note is the much improved fit, for all brands, and the entire market. The only obvious deviating brand would be the smallest brand, now quite clearly above the line. The second largest brand perhaps would come close to being considered a deviation (perhaps now particularly only because of the overall excellent fit for all other brands), though the degree of its deviation from the DJ line is more or less the same degree as it was in the initial, brand-fit method. The largest brand now has almost perfect alignment between its observed and theoretical points. The extent of the deviations for the mid-size brands are much the same as they were in the better-fitting market-fit method. And the second of the smallest brands is now also almost perfectly aligned to its theoretical point. Once again, the MSED value decreases, evincing the improved overall fit: 0.36 (c.f. 0.46 and and 0.62).

Category 2: Australian Credit Cards, Fly Buys Members

1 - Brand-size Weighted S Method

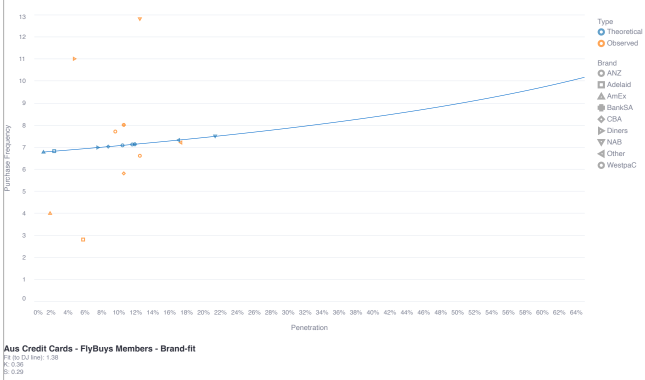

2 - Market Fit S Method

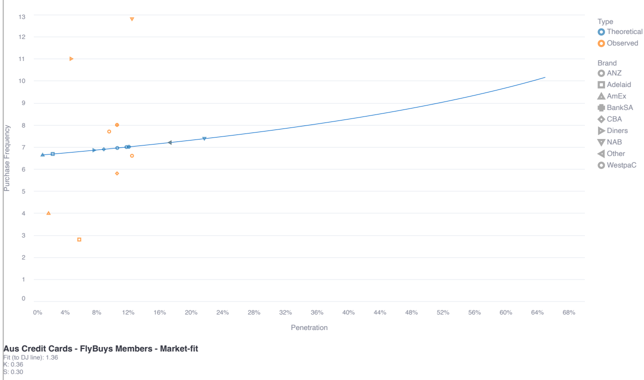

3 - Joint K&S Solution Method

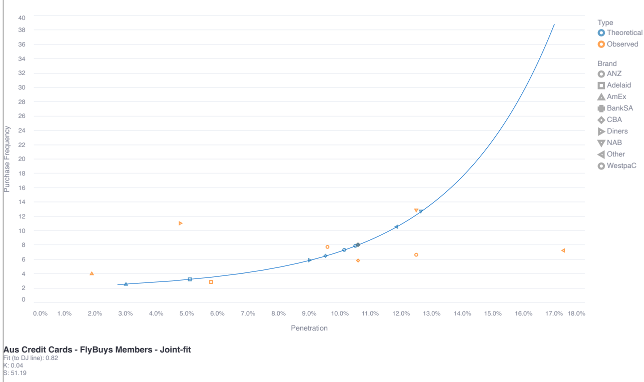

The second category is different to the first in that the fits derived from both the traditional, brand-fit and market-fit methods are both very poor. A straight line of best fit drawn through the observed data points for the brands would be something approaching vertical, going from something like 7 o'clock to 1 o'clock on a clock face. Whereas the line of the DJ curve for the category with these parameters would bisect that near-vertical line, making an almost cross-like pattern; with the two lines almost rotated 90˚ to one another. Hence, it would appear that there is a very poor fit, with only one brand, the brand with the largest observed penetration, being nearly perfectly aligned with its theoretical point. Once again, the fit being more closely 'tuned' to the largest brand. With the joint-fit method, however, the re-scaling (as was evident in the first category) occurs again, and is even more pronounced in this category. Now the DJ curve more resembles a fit-line passing through a series of data points (rather than a line more or less at right angles to the series of data points). Whereas the MSED was abysmal previously, at 1.38 and 1.36, the MSED now is much improved, though still not a decidedly good one, at 0.82 (less than 0.5 would appear to be a useful target).

The still less-than-ideal fit is evident from the extreme deviation of the point with the largest penetration (the collective grouping of 'Other' brands), which now shows severe deficit loyalty along with the brand at about 5% Penetration. Of course, the analyst could choose to omit from the analysis the combined grouping of multiple small brands into an aggregate larger brand.

Conclusions

This Technical Note has documented new approaches to fitting the NBD-Dirichlet model. These approaches use the Euclidean Distance of the observed to the model-predicted theoretical data points of Penetration and Purchase Frequency to calculate the mean standardised distance as a measure of fit.

The fit measure is then used in model fitting to either i) optimise the fit by solving for the value of the S parameter, given a previously derived (through 'traditional' methods) value of the K parameter; or ii) optimise the fit by jointly fitting K and S simultaneously.

These new methods of fitting the NBD-Dirichlet model do yield different values for S (or both S and K in a second method) to the value derived from the traditional method. Initial examination, albeit of limited scope, has shown these new approaches do generally result in better fits across the whole market — lower Mean Standardised Euclidean Distance — than the traditional approach. Further, there are circumstances (as illustrated here in this analysis) where there can be materially different interpretations derived from the disparate fit methods. This has implications for analysts and researchers looking to investigate the fit of (often to identify and comment on deviations from) the predictions of the NBD-Dirichlet model to observed, real-world data. A data point classified as a deviation (e.g. excess or deficit penetration) through one method may not remain as being classified as such under a different fitting method.

Reference List

Aldenderfer, M.S. and Blashfield, R.K., 1984. A review of clustering methods. *Cluster analysis*, pp.33-61.

Bass, F.M., 1993. The future of research in marketing: marketing science. *Journal of Marketing Research*, *30*(1), pp.1-6.

Bass, F.M., 1995. Empirical generalizations and marketing science: A personal view. *Marketing Science*, *14*(3_supplement), pp.G6-G19.

Bound, J., 2009. User's Guide to DIRICHLET. *Marketing Bulletin*, *20*.

Chen, F. and Chen, M.F., 2008. The Dirichlet package.

Driesener, C., 2005. *Empirical generalisations in the parameter values of the NBD-Dirichlet Model* (Doctoral dissertation, PhD Dissertation submitted at the Ehrenberg Bass Institute of Marketing Science, University of South Australia).

Driesener, C., Banelis, M. and Rungie, C., 2017. If the model fits, use it: Methods and benchmarks for evaluating NBD-Dirichlet goodness-of-fit. *Australasian Marketing Journal*, *25*(4), pp.288-293.

Ehrenberg, A.S., 1959. The pattern of consumer purchases. Journal of the Royal Statistical Society Series C: Applied Statistics, 8(1), pp.26-41.

Ehrenberg A.S.C. (1988), *Repeat Buying: Facts, Theory, and Applications.* London: Charles Griffin and Co.; New York: Oxford University Press.

Ehrenberg, A.S., Uncles, M.D. and Goodhardt, G.J., 2004. Understanding brand performance measures: using Dirichlet benchmarks. *Journal of Business research*, *57*(12), pp.1307-1325.

Goodhardt, G.J., Ehrenberg, A.S. and Chatfield, C., 1984. The Dirichlet: A comprehensive model of buying behaviour. *Journal of the Royal Statistical Society. Series A (General)*, pp.621-655.

Greenacre, M. J. (1984). *Theory and applications of correspondence analysis*. London: Academic Press Inc.

Habel, C. and Lockshin, L., 2013. Realizing the value of extensive replication: A theoretically robust portrayal of double jeopardy. *Journal of Business Research*, *66*(9), pp.1448-1456.

Kalwani, M.U., Meyer, R.J. and Morrison, D.G., 1994. Benchmarks for discrete choice models. *Journal of Marketing Research*, *31*(1), pp.65-75.

Kearns, Z., 2010. Dirichlet. xls. Palmerston North. Massey University.

Pare, E.G., 2000. Loving R O. Hill I. Descriptive Geometry Metric.

Rungie, C. and Goodhardt, G., 2004. Calculation of Theoretical Brand Performance Measures from the Parameters of the Dirichlet Model. *Marketing Bulletin*, *15*.

Rungie, C.M., Laurent, G. and Stern, P., 2003. Modeling Long-Run Dynamic Markets. *Adelaide: School of Marketing, University of South Australia*.

Sharp, B.M., 1997. *Evaluating the impact of a loyalty program on brand loyalty* (Doctoral dissertation, PhD Dissertation submitted at the University of Adelaide).

Uncles, M.D., 1989. Buyer: Buyer Behaviour Software: a Manual for Version 89.1. London Business School.

Uncles, M., Ehrenberg, A. and Hammond, K., 1995. Patterns of buyer behavior: Regularities, models, and extensions. *Marketing Science*, *14*(3_supplement), pp.G71-G78.