The Temporal Stability of a Stochastic Model

Malcolm Wright, Ehrenberg-Bass Institute, University of South Australia,
malcolm.wright@marketingscience.info
Lara Stocchi, Ehrenberg-Bass Institute, University of South Australia,
lara.stocchi@marketingscience.info

Abstract

We report on the stability, over time, of the parameters of the NBD-Dirichlet model. This model is widely used and is recommended as a stationary market benchmark against which the dynamic effects of marketing interventions can be assessed. Yet there is little work assessing how sensitive the model parameters are to the minor fluctuations present in most markets most of the time. Using four data-sets, we assess the stability of the model by fitting it repeatedly to monthly data. The results show the model parameters are generally stable, except that the distribution of heavy and light buyers varies quite markedly in one category (toothpaste), as the ratio of light to heavy buyers occasionally changes in this market.

Keywords: NBD-Dirichlet Model, Temporal Stability, FMCG, Brand Loyalty

The Temporal Stability of a Stochastic Model

Introduction

The NBD-Dirichlet model of purchase incidence and brand choice is one of the most well-established empirical generalisations in marketing. It is widely used by organisations such as Bases, Procter and Gamble, and Kraft. It provides benchmarks against which consumer behaviour and marketing effects can be interpreted. These benchmarks include the predicted distribution of heavy and light buyers, the expected duplication of purchases with competitive brands, and a variety of brand performance metrics such as the number of buyers, average purchase frequency, share of category requirements and the proportion of solely loyal buyers.

Both the NBD and Dirichlet components of the model also have wider applications; for example they may be adopted as prior distributions in Hierarchical Bayes analysis. See Jen, Chou and Allenby (2003) for an example of an NBD prior distribution, and Chaing, Chib and Narasimhan (1999) for an example of a Dirichlet prior distribution.

The model assumes that the markets analysed are stationary, and Ehrenberg (1988) justifies this assumption by claiming that as a matter of fact markets *are* broadly very stable, with most deviations failing to persist over multiple analysis periods. However, there are some systematic deviations, such as a variance discrepancy in the NBD part of the model, and a slight under-prediction of solely loyal buyers in the Dirichlet. Others have also identified systematic deviations from the NBD-Dirichlet model, as in Fader and Schmittlein's (1993) seminal paper on excess loyalty for high share brands. Generally, these deviations are small, and are known boundary conditions for the model's estimates of brand performance metrics.

While there has been a moderate amount of empirical work on the deviations for brand performance metrics, there has been surprisingly little research into the stability of the underlying model parameters. One partial exception is Wright and Sharp (2001), who examined changes in model parameters during a new product launch. Driesener (2005) has also examined parameter differences between categories. However, we could find no work on parameter stability in mature markets. This is a curious omission in the methodological work underlying this important model, especially given the fact that some probabilistic models are known to be subject to parameter instability – such as exponential gamma models of new product trial (Fader, Hardie and Zeithammer, 2003).

Therefore, in this paper we report an exploratory study of parameter stability for the NBD-Dirichlet model with up to three years of data for four FMCG products. We are interested in whether the model parameters vary over time or fluctuate in response to market events. This work is an important part of the checking and validation process supporting the use of the NBD-Dirichlet model: in industry applications; in analysis of consumer behaviour and brand performance; and supporting its use as a prior distribution in Hierarchical Bayes analysis.

Data and Method

We used data from the Taylor Nelson Sofres (TNS) UK panel, accessed and processed using the TNS database system Powerview. This provides aggregate, rather than individual outputs, which in turn required the Dirichlet Model to be fitted with the method of Means and Zeros, as originally outlined by Goodhardt and colleagues (1984), rather than with maximum

likelihood estimation method. We used the well-known Excel software for Dirichlet analysis by Kearns to undertake this estimation. We examined 4 FMCG categories: breakfast cereals confectionary, toothpaste and deodorant. In each case we analysed the top brands accounting for most of the market.

Following standard practice for means and zeroes estimation, we input the observed sample mean proportion of non-zero buyers (penetration b, referred as capital B for the category) and the average rate of purchase (w, referred as capital W for the category) for both a variety of single brands and also for the entire market, see also (Morrison, 1969), (Goodhardt, Ehrenberg and Chatfield, 1984), (Morrison and Schmittlein, 1988). This enabled estimation of the parameters for the mathematical distributions compounding the Dirichlet Model (i.e. the Negative Binomial Distribution and Dirichlet Distribution) and the generation of a wide range of theoretical brand performance metrics. The means and zeroes methods does not require all brands to be included, as category purchase incidence is estimated from category statistics, while brand switching relies on estimates from one or more (weighted) brands rather than required data from all brands.

We fitted the model repeatedly to each month of data for each category. This provided a time series of parameter values that allowed an assessment of temporal stability, both informally through plots of the time series and more formally by considering the ratio of the standard deviations to the means of the parameters (that is, the coefficient of variation).

Results

Figure one plots the stability of the M parameter – which reflects normalized sales. This is the parameter in which we would expect most variation, as it is responsive to seasonal and promotion effects.

Figure 1: Time Series of Monthly NBD-Dirichlet M Parameter

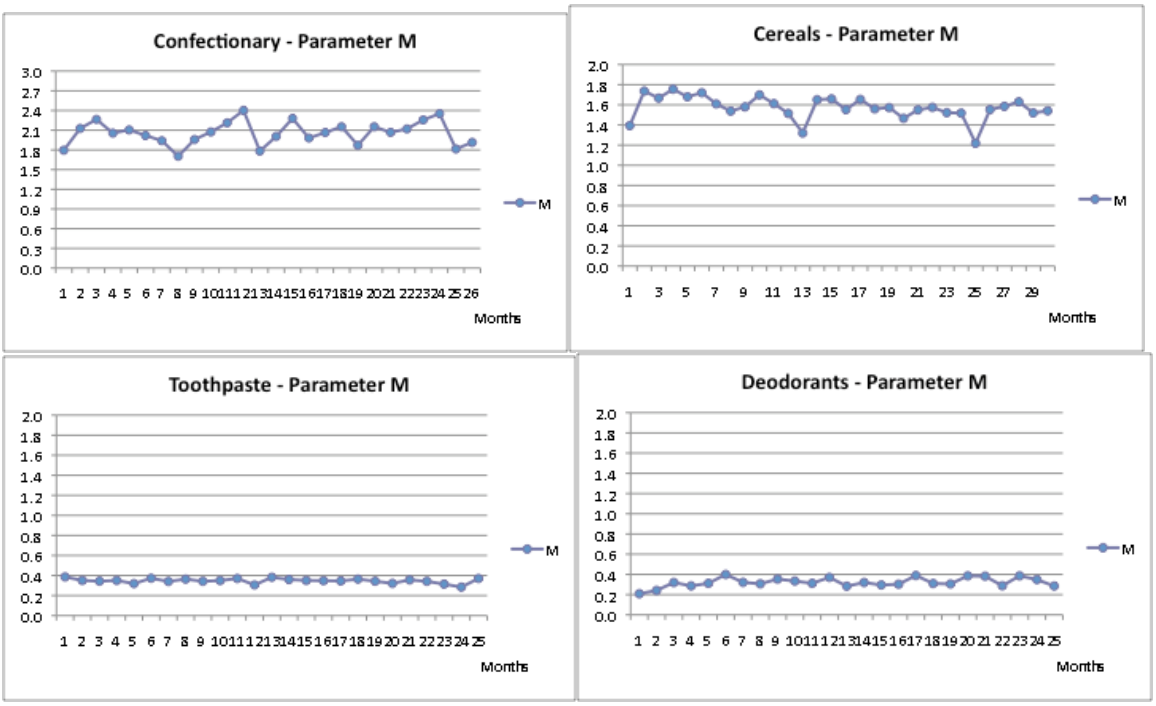

As can be clearly seen, there was relatively little period-to-period variation in M. The coefficient of variation ranged from .07 (for cereals and toothpaste) to .15 (for deodorants) and averaged just .10. Therefore, temporal variation is small relative to mean values for M.

Figure 2 examines the K parameter of the NBD component of the model. This parameter captures the distribution of mean buying rates across individuals – in effect the mixture of light and heavy buyers. Here again, we see considerable stability except in a single category: toothpaste. The coefficient of variation averages just .23 over all four categories – however for three of the categories is is quite low (.05 to .10) whereas for toothpaste it is .70, reflecting wild swings in K. This shows that from time to time the distribution of heavy and light buyers changes rapidly. Increases in K are typically associated with increases in the percentage sales importance of light buyers, so the spikes in K represent relatively more light buying and less heavy buying.

Figure 2: Time Series of Monthly NBD-Dirichlet K Parameter

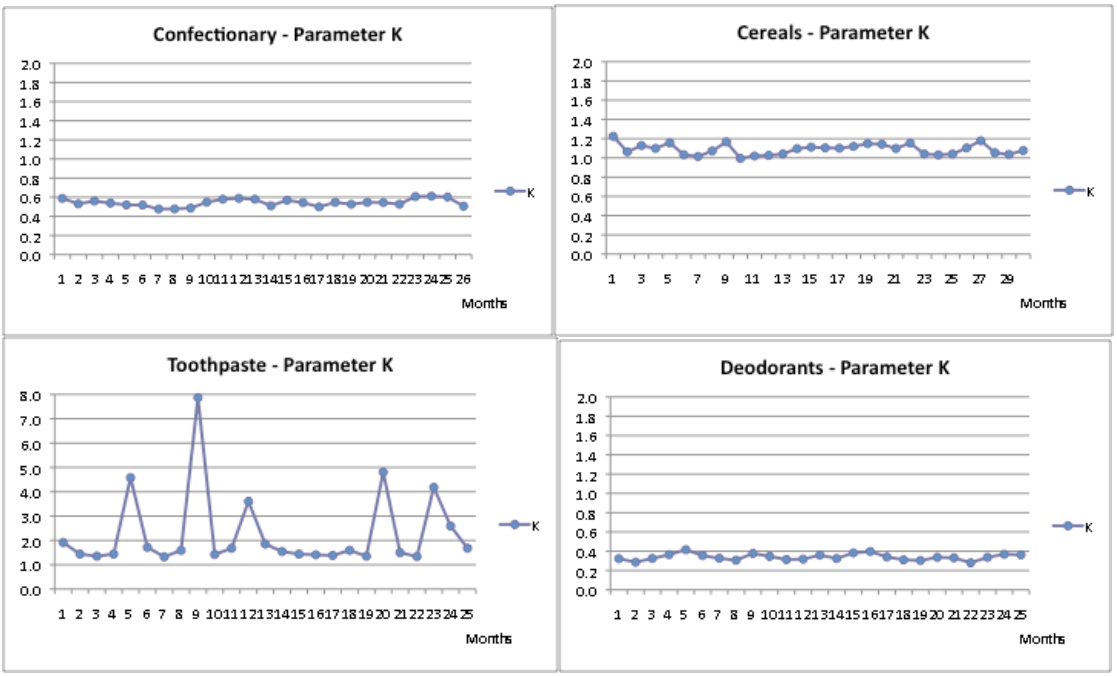

To better understand the variations in K we would ideally examine individual purchase records and the detail of marketing interventions at the time. Such analysis is, regrettably, beyond the scope of this particular paper. However, logically, the spikes must be due either to light buyers coming into the market, or heavy buyers dropping out of the market. The first possibility would be consistent with a promotion at the time of the spike, bringing light buyers into the market. The second would be consistent with a promotion just before the spike, leading to purchase acceleration and stockpiling by heavy buyers. Whichever the case, we can hypothesise that toothpastes are a promotion-sensitive category, although this should be checked with further data and analysis. Nonetheless, this figure does demonstrate how analysis of the NBD-Dirichlet parameters can yield some interesting findings.

Figure 3 turns from category stability to brand stability. It reports the Phi parameter, which is a transformation of the S parameter of the Dirichlet component of the model. In effect, Phi measures the degree of loyalty within the category. Although this is primarily a category measure, it will sum and reflect the individual brand's levels of loyalty. If a brand has a substantial increase in loyalty, this will be reflected in the Phi parameter.

In this case, we again see substantial temporal stability. The average cofficient of variation across the four data-sets is just .19, ranging from a low of .09 to a high of .26. This shows that not much changes in brand loyalty over the time periods and categories anlayzed.

Figure 3: Time Series of Monthly NBD-Dirichlet Phi Parameter

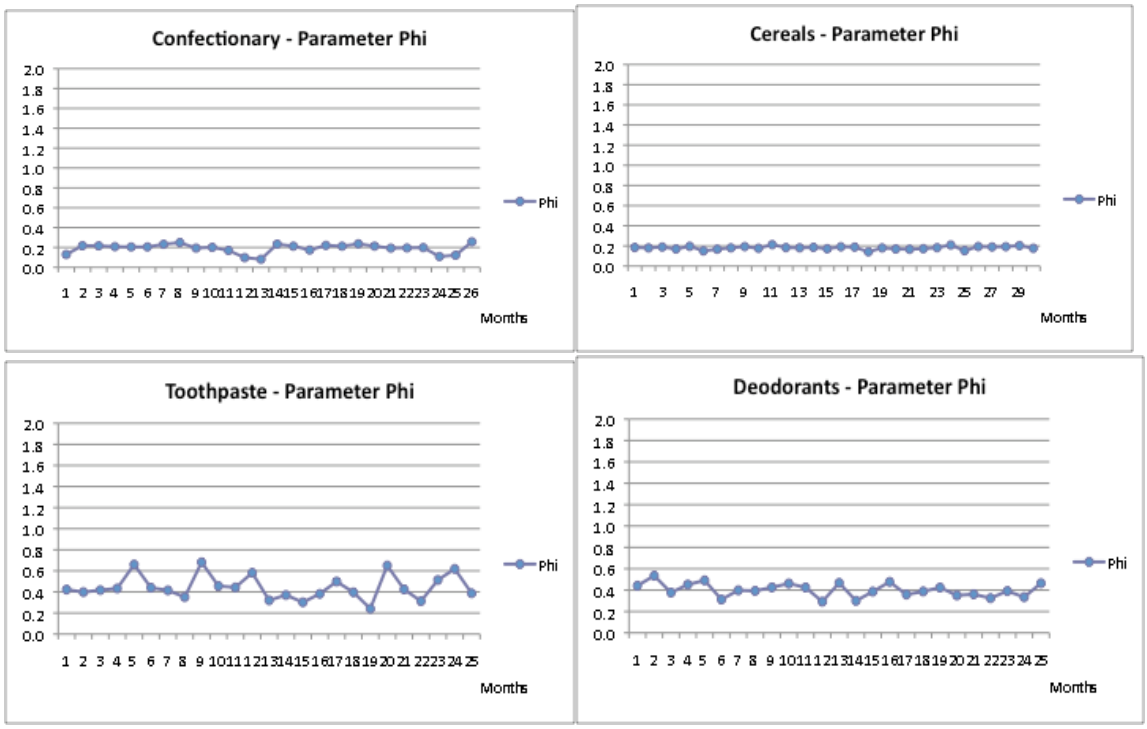

Conclusion and Future Research

This analysis suggests that the NBD-Dirichlet model is stable, and is robust to minor seasonal variations. With one exception (toothpaste) the parameters show little variation over a long period of time. Overall, the NBD-Dirichlet model passes the test; just as Ehrenberg and Goodhardt claimed (1988), markets are largely stable. This examination of four FMCG markets shows, for the first time, that this stability typically extends to the parameters of the model, albeit with a single exception. However, even with this exception included, the model parameters are much more stable than is found in the aforementioned exponential-gamma models of first purchases for new products. So performance remains relatively good.

What can we conclude about the unusual results for the K parameter for toothpaste? As mentioned earlier, as K goes up (as we see in Figure 2) light buyers are responsible for a greater proportion of sales. Further, the large variations in Figure 2 are in periods 5, 9, 12, 20

and 23, and during these periods we tend to slight drops in M in Figure 1 and increases in Phi in Figure 3.

To put it another way, in the time periods that light buyers account for a greater proportion of sales, sales also tend to drop slightly, and brand loyalty increases slightly. This is suggests that rather than getting more light buyers into the market, we are actually see some heavy buyers temporarily out of the market, resulting in lower demand and less brand switching (light buyers are generally more loyal, out of necessity.) The reason for these patterns is not clear at this stage. More research is needed.

We note in passing that the K parameter does consistently revert back to the baseline level, as Ehrenberg would have expected. However, the regularity of the fluctuations is not at all what he had in mind.

At this stage we do not propose more analysis of this exploratory and descriptive results. They are interesting in their own right – hence the conference presentation. However, they are also part of a much more comprehensive research program. We have work underway to examine bias, as well as stability, to apply a suite of measures of model fit, and to consider whether there are aggregation effects in bias and stability. The results of this more comprehensive work will be reported in due course, and from these we may be able to develop greater insights into the unusual patterns in the toothpaste category.

References

Chaing, J., Chib, S., & Narasimhan, C., 1999. Markov chain Monte Carlo and models of consideration set and parameter heterogeneity. Journal of Econometrics, 89, 223–248.

Ehrenberg, A.S.C., 1988. Repeat-buying: facts, theory and applications, Oxford University Press, London.

Fader, P.S., Hardie, B.G.S. and Zeithammer, R., 2003. Forecasting New Product Trial in a Controlled Test Market Environment. Journal of Forecasting 22, 391-410.

Fader, P.S. and Schmittlein, D.C., 1993. Excess Behavioral Loyalty for High-Share Brands: Deviations from the Dirichlet Model for Repeat Purchasing. Journal of Marketing Research 30 (November), 478-493.

Goodhardt, G.J., Ehrenberg, A.S.C. and Chatfield, C., 1984. The Dirichlet: A comprehensive model of buying behaviour. Journal of the Royal Statistical Society 147 (5), 621-655.

Jen, L., Chou, C.H., & Allenby, G. M. , 2003. A Bayesian approach to modeling purchase frequency. Marketing Letters, 14, 5–20.

Morrison, D.G., 1969. A series approximation for negative binomial parameter estimation. Journal of Marketing Research VI (August), 355-356.

Morrison, D.G. and Schmittlein, D.C., 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), 145-159.

Wright, M. and Sharp, A., 2001. The Effect of a New Brand Entrant on a Market. Journal of Empirical Generalisations in Marketing Science 6, 15-29.