A SYSTEMATIC INVESTIGATION OF THE ESTIMATION OF THE DIRICHLET MODEL

A thesis presented in partial fulfilment
of the requirements for the degree of Doctor of Philosophy
in Marketing at Massey University

Zane Kearns

1999

ABSTRACT

The NBD/Dirichlet is a stochastic model of purchase incidence and brand choice which parsimoniously integrates a wide range of well-established empirical regularities in fast moving consumer goods markets. More recently this work has been extended into other areas such as the prescribing of pharmaceuticals, (Stern 1994); airline aviation fuel contracts, (Uncles and Ehrenberg 1990a); and the visiting of retail store chains, (Uncles and Ehrenberg 1990b).

By combining the stochastic assumptions of the model, namely Poisson purchasing of products, with mean rate distributed gamma across the population, and brand choice represented by multinomial probabilities distributed Dirichlet across consumers; a number of aspects of the aggregate behaviour of consumers can be successfully predicted successfully.

This thesis examines the estimation issues in the Dirichlet model, specifically, the central Dirichlet parameter *S* used to represent heterogenity in brand choice.

ACKNOWLEDGEMENTS

In the first instance I would like to thank Professors A S C Ehrenberg, G J Goodhardt, and C Chatfield for their development of what has become known as the Dirichlet model. It is difficult to discern the unique contributions of these researchers, but overall their efforts represent a serious but often unrecognised contribution to marketing and the study of consumer behaviour.

I would also like to thank Professor Ehrenberg more personally for the time taken to answer my queries about the model, and being overly kind in not pointing out how silly some of my questions were, or that I should really have known the answer.

Of course the greatest burden of gratitude goes to my supervisor Associate Professor A C Lewis. His guidance is evident throughout this thesis, and without it, the thesis would be undeniably poorer. Thanks must also go to Dr Greg Arnold who lent valuable advice.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

CHAPTER ONE

INTRODUCTION

1.1 Background

The Dirichlet model is an empirical generalisation that describes certain patterns and aspects of the aggregate behaviour of the buyers of competitive brands. These patterns and aspects refer to the purchasing of a brand in a period; purchasing of a brand in two periods; and the distribution of a consumer's purchases in a product category.

Perhaps the most widely known finding is that these patterns relate to the sales level or market shares of brands, especially in stationary markets, Ehrenberg (1995a). Ehrenberg goes on to say that in general, many observed patterns of buying do not depend on the brand or product itself, nor on what else the buyers of the brand buy, nor on external factors such as advertising, pricing, and distribution. Instead the patterns of behaviour can be explained simply by factors such as how many people buy the brand and how often.

Ehrenberg's contention is thus summarised:

> Of the thousand and one variables which might affect buyer behaviour, it is found that nine hundred and ninety-nine usually do not matter. Many aspects of buyer behaviour can be predicted simply from the penetration and the average purchase frequency of the item, and even these two variables are interrelated (Ehrenberg, 1988, pg. ii).

The many regularities of buyer behaviour are remarkable. The Dirichlet model integrates the reported regularities, and predicts many aggregate brand performance measures. These measures are the distribution of purchases for a brand, the proportion of a brand's buyers buying that brand only, and the proportion of people purchasing a brand, given that they have previously purchased that brand. When these predictions are compared with observed figures, Ehrenberg claims that it is not unreasonable to expect to obtain correlations in the order of 0.9 and sometimes much higher, (Ehrenberg 1975, Ehrenberg and Bound 1993).

Applications of the theory are illustrated in Ehrenberg and Uncles (1995). For example, the theory can be used to provide norms for examining brand performance, or diagnostic information for the "health" of a brand. In addition, the Dirichlet model can provide interpretative norms for evaluating situations where some trend in sales has occurred, say after a promotion or advertising scheme. Ehrenberg also claims that the Dirichlet model provides valuable insights into the nature and implications of brand-loyalty (e.g., Ehrenberg and Uncles 1995; Ehrenberg and Uncles 1999).

Goodhardt, Ehrenberg and Chatfield (1984), summarise the situation by stating that the Dirichlet model makes explicit that there are simple, general and rather precise regularities in a substantial area of human behaviour where this has not always been expected. In setting the context for this particular approach to the modelling of consumer behaviour viz. the largely explanatory models of consumer behaviour, Ehrenberg (1988) claims that it describes how consumers behave, rather than why, and takes into account only those factors necessary for an adequate description.

Ehrenberg develops this when he states:

> Before one can explain the individual consumer's decision processes and behaviour, one needs to know and understand the overt behaviour that has to be explained - what generalisable regularities there are and what apparent inconsistencies, And knowing the factors from which one can successfully predict consumer behaviour (and especially those factors which do not matter in this respect) does in fact provide major insights into its nature, (Ehrenberg 1988, pg. viii).

1.2 Objectives of the Thesis

The important performance measures predicted by the Dirichlet model are the proportion buying once, twice and so on; the proportion who only buy that brand, and their purchasing rate; the rate of category purchasing; which other brands are bought; and period to period repeat buying of a brand. In addition, these predictions can be made for different length time periods.

The predictions listed above rely on estimates of a number of parameters. These are: the proportion buying the category, and the average rate of buying of buyers; the proportion buying a brand, and the average rate of buying of the brand buyers; K, the heterogeneity parameter for the distribution of purchases at the category level; and S, the brand choice heterogeneity parameter.

Numerous developments of the model have taken place since the first model was suggested in Ehrenberg (1959). Ehrenberg summarises further work on the current model that is desirable:

(i) Prediction errors require investigation. Ehrenberg claims the errors are mostly small and may be of little marketing importance.

(ii) It would be helpful to develop closed-form approximations to the Dirichlet.

(iii) The lack of constancy of the parameter S should be investigated. As it stands, different values are obtained from different length time periods, and from different samples in the same market.

(iv) The issue of whether the two basic "diversity" parameters K and S are product class characterisics, and whether they are the same for different demographic sub-groups, different store-groups, different countries and so on, needs to be investigated.

(v) The meaning and interpretation of the S parameter should be clarified. Ehrenberg (1995b) goes on to ask "how far is the variation in S mere sampling? What other explanatory factors are involved? What effects does the variation in S have on the predictions of the model?", (Ehrenberg 1988, pg. 130).

From this summary the following thesis objectives have been formulated.

(1) To assess the effect of sampling error on the estimates of S.

(2) To assess the impact of different values of S on predictions obtained from the Dirichlet model.

1.3 Outline of Chapters

Chapter Two describes the NBD/Dirichlet model. The chapter begins by summarising the main features of the model and goes on to illustrate ways in which the predictions of the model can be used.

Chapter Three examines the assumptions underlying the model. The purpose of this chapter is to review the justification for the model, and possible alternative structures which might lead to better predictions.

Chapter Four reviews the literature relating to the estimation of the Dirichlet model. The main aim of the chapter is to describe a range of methods that have been proposed.

Chapter Five details the specific objectives of this thesis, and then describes the methodology used to meet the objectives.

Chapter Six sets out the findings of this thesis, and Chapter Seven discusses the implications of the research findings.

CHAPTER TWO

STOCHASTIC MODELLING OF CONSUMER BEHAVIOUR

2.1 Introduction

This chapter gives an overview of the development of the Dirichlet model first as part of the marketing literature concerning brand-loyalty, and then as a contribution to the stochastic modelling tradition in marketing. The chapter concludes by examining the fit of the model's predictions.

2.2 Loyalty

There is an extensive literature on the topic of brand-loyalty. The word loyalty is, according to Uncles and Laurent (1997), used widely in marketing literature and the marketing plans of practitioners. An extensive review of the brand-loyalty literature published nearly twenty years ago found there were more than 300 published studies on brand-loyalty (Jacoby and Chestnut 1978). Since that review there seems to have been little decline in interest in the concept of brand-loyalty.

For example, Dekimpe, Steenkamp, Mellens and Abeele (1997) state that a critical issue for the continued success of a firm is its capability to retain its current customers and make them loyal to its brands. Cited justification is that it has been found that the cost of attracting new customers is six times as high as the costs of retaining old ones (Rosenberg and Czepiel 1983), loyal customers are said to be less price sensitive (Krishnamurthi and Raj 1991), and the presence of a loyal customer base provides the firm with valuable time to respond to competitive actions (Aaker 1991). Furthermore, according to Fournier and Yao (1997), "after a long hiatus on brand-building activities, marketers have once again placed the development of consumer brand loyalties at the heart of their business plans". Sharp and Sharp (1997) also report a rise in the use of loyalty programs by marketers.

Much has also been written about the apparent decline in brand-loyalty. Dekimpe, et al. (1997) present examples of some of these claims. For example, brand- loyalty is often said to be being replaced by price-loyalty (Anon., Discount Merchandiser, 1993). The increasing fragmentation of the market (Anon., Marketing 1993), and the growing popularity of cheaper regional and private-label brands (Anon., Brandweek 1993) have also been cited as reasons for an apparent decline in brand-loyalty in recent years. Pfouts (1994) claims that the diminishing brand-loyalty, especially in food items, is one of the most striking revolutions in recent years, while an article in Industry Week (Anon., Industry Week 1993) also claims that brand-loyalty is a thing of the past. Earlier, Dodson, Tybout and Sternthal (1978) and Strang (1975) argued that the reliance of many national brands on price promotions will be harmful to their long term health (meaning loyalty).

In contrast with these various opinions about brand-loyalty are the opinions informed by the stream of research culminating in the Dirichlet theory as described, for example, by Ehrenberg (1988). As Uncles and Laurent (1997) point out, behavioural patterns of loyalty are reasonably well understood and integrated in the Dirichlet theory. Fader and Schmittlein (1993) also note that the Dirichlet theory has been used by numerous researchers to explain choice patterns made by heterogenous decision makers, and that the Dirichlet offers a robust, parsimonious method to summarise and predict repeated choices. It also is said that the Dirichlet model offers a variety of diagnostic statistics with useful managerial implications (Fader and Schmittlein 1993; Goodhardt et al. 1984).

Bhattacharya is a strong advocate for the Dirichlet model when he states:

> Researchers such as Ehrenberg have repeatedly established that simple parameters such as penetration, purchase frequency and market share are capable of predicting many other aspects of consumer behaviour including behavioural brand-loyalty. Such prediction is easily accomplished using the Dirichlet model of consumer purchasing behaviour. The predictive capabilities of the Dirichlet are well known, (Bhattacharya 1997, pg. 421).

Despite the extensive empirical literature concerning the Dirichlet model most of the work on brand-loyalty pays little attention to the predictive capabilities of the Dirichlet model. In a recent debate concerning brand equity Ehrenberg (1995) challenged critics to demonstrate consumer based measures of brand equity which were not functions of the brand size, something that they were unable to do, except noting that such measures remain theoretical possibilities.

Bhattacharya, Fader, Lodish and De Sarbo (1996) note that many authors have discussed the shortcomings of loyalty measures based solely on purchase histories (see for example Lattin 1990; Little and Anderson 1994; Ortmeyer, Lattin and Montgomery 1991; Russell 1994; Srinivasan and Kibarian, 1989). One consistent reason for the criticism may be that behavioural demonstrations of loyalty have often been criticised as only one aspect of loyalty. For example, Dekimpe et al. (1997) cite the importance of an attitudinal component. Furthermore, Ambler (1995) and Baldinger and Rubinson (1996) criticise the aggregative (behavioural) measures employed as not providing a true indication or managerially interesting perspective of "loyalty".

However, as Colombo and Morrison (1989) note, behavioural data refer to what consumers do, and therefore should at the very least be used as a benchmark or test of convergent validity of any other measure. Recent work by Barnard and Ehrenberg (1990), Castleberry, Barnard, Barwise, Ehrenberg and Dall'Olmo Riley (1993), and Dall'Olmo Riley et al. (1997) also raises the issue that there are some inherent difficulties with attitudinal measures of brand-loyalty conventionally used. Specifically, this work is consistent with the hypothesis that many measures of attitudinal-loyalty vary with current behaviour rather than reflecting stable evaluations of brands which might be thought to cause behaviour.

Ehrenberg, Uncles and Hammond (1995) also point out that the Dirichlet prediction of similar behavioural loyalty measures for similar sized brands also has implications for theories about consumers' attitudes about brands. Thus, brand images should generally not be expected to differ markedly between similar brands, except for market share effects and occasionally some highly "descriptive" attributes.

At the same time that Ehrenberg was making his repeated claims about regularities in buyer behaviour, there was an exponential growth in data. The main data applicable to investigations of consumer behaviour are records of what people buy, day-by-day or week-by-week. Such records are primarily obtained from so called "consumer panels". These are market research operations where all purchases in a specified range of product-classes are continuously measured for a sample of potential customers. The measurement procedures used can vary in detail. Early attempts required consumers to keep a diary of their purchases or partake in regular interviews to recall purchases. More recently, most panels have been updated to scanner panels, where consumer purchases are recorded either at the point of purchase or in the household.

2.3 Stochastic Modelling of Consumer Behaviour

Bass, Givon, Kalwani, Reibstein and Wright (1984) state "that is now well established that the fundamental nature of consumer brand choice behaviour for frequently purchased, low-priced products is one of frequent switching among brands", pg. 267. Bass et al. go on to summarise that the application of stochastic models of consumer purchasing behaviour played an important role in providing parsimonious descriptions of such purchasing behaviour

For example, Bass, Pessemier and Lehmann (1972) found in a study of soft-drink buying, that subjects bought the brand they most preferred 55 percent of the time, only 45 percent bought the brand that they bought last time, and only 41 percent bought the brand that they had the most favourable attitudes to. Bass, Jeuland and Wright (1976) also summarise evidence that indicates that the brand choices of consumers do not remain constant through successive trials even though the stated preferences of subjects are unchanging. Much work has attempted to integrate such findings into theories of behaviour.

There are alternative ways to formulate stochastic models of consumer behaviour. Bass et al. (1976) point out that the various models in the main differ in assumptions about the underlying stochastic process which results in the observed behaviour. Massy et al. (1970); Montgomery and Ryan (1974); Lilien, Kotler and Moorthy (1992); and Dalal, Lee and Sabavala (1984) give summaries of the development of stochastic models in marketing.

The application of a stochastic model to purchasing behaviour usually specifies some probability law for the purchase patterns of an individual customer, and allows the parameters of the specified customer model vary over the population according to some probability distribution. For example, within the Dirichlet model under consideration here, it is assumed that purchases within a time period for an individual consumer are consistent with a Poisson process.

Stochastic models of consumer behaviour also have traditionally been of two main types; purchase incidence models and brand choice models (Blattberg and Sen, 1976). This just means that the stochastic modeller has to make decisions about the incorporation of appropriate representations of the processes required to generate the observed behaviour that is to be modelled. Bass et al. Wright (1976) point out that distinctions can also be drawn between dynamic (where the probabilities specified in the model are expected to change over time), and steady state versions of the theories.

According to Ehrenberg and Uncles (1995) most published works on mathematical models of buying behaviour have been mainly concerned with how markets change rather than describing the steady state as is done in Dirichlet analyses of markets.

Ehrenberg and Uncles (1995) recognise alternatives to the Dirichlet model. They claim the main difference, is that the alternatives are do not offer as wide a range of generalisable empirical findings. Such a position led Uncles, Ehrenberg and Hammond (1995) to claim that the Dirichlet model may be the best-known example of an empirical generalisation in marketing, with the possible exception of the Bass diffusion model.

In a wider context, Ehrenberg et al. (1990) and Bass (1974) note that theorists of consumer behaviour (e.g., Howard and Sheth 1969; Engel, Blackwell and Miniard 1995, Cooper and Nakanishi 1988; Peter and Olson 1995), must acknowledge, or at least allow implicitly, that buyers of smaller brands tend to be less loyal. That is, that any theory that involves consumers, such as how pricing or advertising might work, must at least be consistent with the patterns described by the Dirichlet model.

2.4 The Dirichlet model of Consumer Behaviour

The Dirichlet comes as a culmination of forty years of routine analyses of consumer panel data, and the regularities noted, (Uncles, Ehrenberg and Hammond; 1995). Seminal contributions outlining the development of the model are provided by Chatfield, Ehrenberg and Goodhardt (1966) and Goodhardt et al. (1984), and the summary provided in Ehrenberg (1988).

2.4.1 Aggregative Considerations

The application of the Dirichlet model requires that a market be stationary and unpartitioned (sometimes referred to as being unsegmented). The specific meanings of these assumptions are discussed below. A second set of assumptions are made about the probability distributions for purchase incidence and brand choice and these are reviewed in Chapter Three.

Stationarity

Consumer purchasing behaviour is dubbed "stationary" when the aggregate level of purchasing remains the same from one time period of a given length to the next period of the same length. Individual purchasing behaviour will, however, still vary from one such period to the next.

Ehrenberg (1988) characteristically states that stationarity is the rule rather than the exception when considering real market data, or as Bass and Pilon (1980) summarise - market share is in long run equilibrium, but is temporarily perturbed by marketing activities . However, Ehrenberg does concede in practice it is unusual to observe complete stationarity in a given data-set, and most cases studied are one of near-stationarity. Ehrenberg goes on to clarify this by adopting an arbitrary heuristic of not more than ± 10 percent in the mean per-capita rate of buying, in the penetration of the brand, or in the rate of buying per buyer .

Barnard and Ehrenberg (1997) note that the view that in the medium term most markets are near stationary most of the time, is increasingly widely supported in the literature (e.g., Dekimpe and Hanseens, 1995; Lal and Padmanabhan, 1995).

However, the restriction to stationary markets is at odds with many beliefs about markets, and those typified, for example, by Lenk, Rao and Tibrewala (1993) who noted how week-by-week sales of an item can vary markedly. However in a reply to this analysis, Barnard, Ehrenberg, Hammond and Uncles (1994) recommend the use of fairly long base periods (e.g., 4, 8, or 12 weeks) to make good estimates of the parameters. They go on to note that this need not be viewed as being too restrictive since any diagnostic test can be for a period as short or long as required.

Barnard et al. (1994) also make clear the importance of the determining the time period deemed suitable for analysis. The application of the Dirichlet model to a particular data-set requires a careful balancing of the length of the time period under consideration. If it is too short, one is likely to see short term market fluctuations; if it is too long, long term trends (the growth of brands) are unaccounted for. There is little guidance on this point beyond the need to consider periods of time which are some multiple of the average interpurchase time, but clearly long term non-stationarity in the market would question the validity of a stationary-assumed Dirichlet analysis.

Doubts are often expressed about restriction to a stationary or near stationary situation. The question is often raised whether there is much point in examining "stationary" conditions when the aim of much marketing effort is to move from equilibria. However, Ehrenberg (1988) points out that the restriction to near-stationary markets need not be viewed as particularly restrictive. For example, he points out that if one wants to create change, it is important to understand the stationary no-trend situation from which one want to depart. And to evaluate what change has been achieved, one must compare the results with what would have happened in the absence of change.

Barnard et al. (1994) also make an additional defense when they summarise that the empirical evidence related to short-term non-stationarity shows no consistent departures.

Market Partitioning and Segmentation

The patterns modelled by the Dirichlet are said to hold between competitive brands. Brands are said to be competitive if the proportion of purchases devoted to any particular brand is independent of the way the remaining purchases are distributed between other brands. If these probabilities are not independent then some brands compete more amongst themselves than with other brands, and this is often evidence of a partition within the market.

A number of examples are typically given of partitioned markets. For example, the coffee market may be partitioned into instant caffeinated, instant de-caffeinated, and ground coffee partitions, and the petrol market into leaded and unleaded. As an illustration, an individual's probability of buying unleaded petrol is not independent of buying leaded petrol. If a wide definition of the market is used (e.g., all coffee rather than just instant coffee), any sub markets will show up as systematic deviations from the main patterns (and ultimately from the theoretical Dirichlet norms), although it is difficult to identify which part of the variance is due to partition specific factors and which is due to brand specific factors. However, as Ehrenberg (1995a) notes, the data can be re-analysed fitting separate models to the apparent partitions.

Hutchinson and Marchant (1998) are critical of the Dirichlet assumption which they says amounts to saying that there are no systematic "brand positioning effects" which find an echo in consumer behaviour. They say that "commonsense rejects this and is right to do so", pg.31. Barnard, Ehrenberg and Scriven (1998) reply that the Dirichlet describes markets which are unpartitioned but says nothing about whether a market is partitioned or not.

2.5 Dirichlet Findings

Within stationary and non-partitioned markets, the Dirichlet model predicts a number of aggregate brand performance measures. These predictions result from stochastic representations of purchase incidence and brand choice at the individual consumer level, and are discussed further in Chapter Three.

The dominant findings are:

(i) That the different measures of brand performance tend to vary together,

(ii) That the different measures correlate highly with market share (i.e., big brands score higher than smaller ones).

According to Uncles, Ehrenberg, and Hammond (1995) major exceptions to these effects and to the more detailed regularities are rare, although elsewhere Ehrenberg notes that there are some small biases in the model, see for example Ehrenberg and Uncles (1999).

Finding (ii) summarises the extensive empirical regularity of double jeopardy, (see for Ehrenberg, Goodhardt and Barwise 1990). Double jeopardy was noted as early as McPhee (1963), and in the application to buyer behaviour, simply states that a small brand is bought by fewer people, and among those fewer people, the brand tends to be less well regarded. This effect can be viewed across a wide range of attributes both behavioural (for example, purchase frequency or period to period repeat buying) and attitudinal (for example, responses to questions about preference, Castleberry et al. 1993).

Ehrenberg et al. (1990) give an explanation of why this occurs.

> For instance, suppose there are just two restaurants in town, one widely known, and the other more obscure. Suppose also that people who know both restaurants regard them as being of equal merit (equal in quality, service, value for money, accessibility, etc). If people were asked which is their favourite restaurant a DJ effect is bound to occur. The reason is that of the many people who know the popular restaurant, most do not know the more obscure one exists and cannot mention it if asked their favourite. In contrast, of the few people who know the obscure restaurant, most also know the popular one. Hence they will "split their vote" - they are equally likely to mention either restaurant as their favourite (or say "undecided") because we have supposed the two restaurants are of equal merit to those who know both, pg 85.

The performance measures predicted by the Dirichlet model (such as the percent of consumers buying a certain brand in a month; the number of purchases per buyer; the proportion buying once, twice etc.; the proportion who only buy that brand, and their purchasing rate; the rate of category purchasing; and which other brands are bought), should according to Ehrenberg be familiar to many marketing analysts and marketing practitioners. Practical applications include providing interpretive norms, both for stationary and non-stationary markets, and prescriptive uses in marketing planning. Goodhardt et al. (1984) for example, provide a number of illustrative applications.

Schmittlein, Beammor and Morrison (1985) and Dalal et al. (1984), note that the most managerially relevant construct that results from this type of modelling are conditional expectations, namely, the expected number of purchases made in a period given that a consumer purchased in a previous period. Specific examples of conditional expectations are highlighted, for example, as the probability of buying brand *x* in a period given that brand *x* was bought in the previous period (repeat buying); the probability of buying *n*, given *n* purchases in a previous period; and the probability of buying brand *x*, given that brand *y* was bought (purchase duplications).

According to Ehrenberg and Uncles (1995), the Dirichlet model imposes a discipline on marketing planning and evaluation, i.e., it should be used to broadly prescribe available options for the marketing planning activities. This, however, is commonly misconceived, (see for example Little and Anderson 1994), when they ask how the Dirichlet patterns in a stationary market can be of interest to marketing managers who are always trying to change that state. In response Ehrenberg and Uncles (1999) point out that the near-ubiquitous Dirichlet regularities provides a context for decision and evaluation, but do not directly specify the courses of action available to marketing management. For example, the regularity that the average rate of buying a brand hardly differs from brand to brand (except for a small double jeopardy effect), suggests that a marketing plan which aims to increase sales by getting existing buyers to buy more (say twice the market average) would therefore be aiming at something unusual or unlikely. Morrison and Silva-Risso (1995) summarise this as "those companies planning for a higher than average w_i *[purchase frequency]* with their small b_i [penetration] niche brand are bucking the odds more than they think", pg G65. ([] added).

Table 2.1 reproduced from Uncles, Ehrenberg and Hammond (1995) gives an example of the typical presentation of results for a product category.

Table 2.1
Brand performance measures for Laundry Detergent
From: Uncles, Ehrenberg, Hammond (1995).
Observed (O) and Theoretical (T); Annual.

1985: IRI	Market Share*	percentage of Households Buying		Purchases per Buyer		Share of Requirements	
		O%	T%	O	T	O%	T%
Any Brand	100%	94	94*	10.0	10.0	100	100
Tide	25	54	58	4.4	4.1	36	35
Wisk	10	29	31	3.4	3.2	26	25
Bold	8	27	25	2.9	3.0	19	24
Era	6	22	21	2.7	2.9	17	23
Cheer	5	19	17	2.7	2.9	22	22
A & H	5	16	17	3.0	2.9	29	22
All	5	19	17	2.5	3.0	19	24
Fab	5	18	17	2.7	2.9	17	22
Oxydol	4	12	12	2.8	2.8	19	21
Solo	2	9	8	2.4	2.7	16	21
Gain	2	9	8	2.4	2.7	15	21
Ajax	2	9	8	2.4	2.7	13	21
Rinso	2	6	6	2.7	2.7	15	20
Dash	1	4	4	2.8	2.7	24	20
Average brand	6*-	18	18	2.8	**2.9**	21	**23**
Correlation		1.00		0.95		0.74	

* denotes used in fitting of the model

This table shows that the fit for brands of Laundry Detergent for three predictions from the Dirichlet model. Figures for the "average" brand show a close correspondence between the averages for the observed and predicted statistics. At the same time the individual brand fits show some typical variation.

2.5.1 The Fit of the Dirichlet model

The illustrations reported in the previous sections, however, do not allow a judgment to be made about whether the application of the Dirichlet model is justified. If the Dirichlet model is to represent a serious contribution to marketing science it needs to provide a good fit to a wide range of data. On this issue, Uncles, Hammond, Ehrenberg and Davis (1994) report an anonymous commentator as remarking that "the Dirichlet model is clearly not tenable in most markets", p 376, and if this contention is true, this would severely limit the applicability of the model.

Morrison and Schmitlein (1988) note that some consideration has to be given to what one measures the fit of the model across. "Fit" may be used to refer to the predictions obtained from the model, or to the applicability of the assumptions incorporated in the model. The justification of the assumptions underlying the model is considered separately in Chapter Three. In the current context, "fit" refers to the fit of the predictions obtained from the model.

Morrison and Schmitlein (1988) caution that some predictions might be better indicators of the quality of the model than other possible predictions. For example, they claim that the aggregate histogram of purchases may not be a very good indicator of the fit of the model, because deviations from the basic assumptions of the model may show up much more clearly in predictions of conditional expectations than in the histogram of purchases aggregated across consumers. More specifically, the histogram of purchases may still look very similar to the negative binomial distribution, but, the predictions of period-to-period repeat buying may be badly biased (sic).

A qualitative justification of the goodness of fit of the NBD/Dirichlet model is given by the many reported successes of the model, and are summarised, for example, in Ehrenberg and Bound (1993):Table Six, or as in Table 2.2 below.

Table 2.2
Conditions under which the Dirichlet patterns are known to occur.
Adapted from Ehrenberg and Uncles (1995).

Food and Drink
Biscuits, Breakfast Cereals, Butter, Canned Vegetables, Cat and Dog Foods, Cigarettes, Coffee, Confectionary, convenience Foods, Cooking Fats and Oil, Flour, Food Drinks, Frozen Foods, Fruit Squashes, Instant Potatoes, Jams and Jellies, Margarine, Peanut Butter Spreads, Processed Cheese, Refrigerated Dough, Sausages, Soft Drinks, Soup, Take-home Beer.

Cleaners and Personal Care Products
Cosmetics, Deodorants, Diapers, Dishwashing liquids, Disinfectant, Fabric Softeners, Feminie Hygiene Products, Hair Sprays, Laundry Detergents, Paper Tissues, Polishes, Shampoos, Soap, Toothpaste.

Automotive, Health and Information/Entertainment
Aviation Fuel, Motors Cars, Motor Oil.
Over-the-counter Medicines, Pharmaceutical Prescriptions.
Comic Strips, Newspapers, Politicians and parties, Radio Presenters, TV Programs.

Distribution Channels
Chains, Individual Stores, Brands within Chains, TV Channels.

Time and Place
1950 to 1997
Britain, Continental Europe, USA, Australia, New Zealand
Demographic sub-groups
Analysis periods from one week to two years.

Exceptions or Partial Exceptions
Restricted Distribution: Private labels: "Demanding" TV programs.
Some submarkets
Special funding (Luxury Cars, Bribes).
Some Politicians

Ehrenberg and Uncles (1999) note that a common criticism is the difficulty in how an "eye ball" judgement of the fit of the model is made, or in other words "how close was close for us". Ehrenberg and Bound (1993) explain:

> I therefore explain: Fit to us depends on not merely establishing a significant (that is probably non-zero) correlation for a single data-set (SSoD). Instead, it turns on the absence of systematic bias across many sets of data (MSoD) and on correlations between the observed and theoretical values generally on the order of 0.9 , pg.98

Reflecting on the fit of the model Ehrenberg has claimed elsewhere that it is not unreasonable to expect to obtain correlations in the order of 0.9 and sometimes much higher in assessing the fit of the Dirichlet, (Ehrenberg 1975; Ehrenberg and Bound 1993), although, it is not clear what the evidence for this correlation is.

Ehrenberg (1995) goes on to state that "we cannot expect too much from a model whose only brand-specific input is each brand's market share," pg. 96. Despite this Ehrenberg and Uncles (1995) claim that any discrepancies are mostly small and explicable in terms of case specific factors. Barnard et al. (1994) also repeat the claim about the efficacy of the model when they state that given the very parsimonious calibration inputs, the Dirichlet must therefore be oversimplified (with innumerable "excluded variables"), and at times biased even in broadly stationary and unpartitioned markets, although in practice they go on to note that the model only fails to fit when the assumptions of stationarity and partitioning are violated.

Ehrenberg (1997) further clarifies "there are large brands and small ones...plus a bit of as yet incoherent wobble", and "any deviations could be due to a variety of other marketing-mix factors (promotions, out of stock, being de-listed, a new flavour, repositioning, a change in sales director or in price, advertising, etc), rather than any long term idiosyncrasy in the brand's still elusive and undefined Equity", pg 23.

However, more recently systematic investigations of the fit of Dirichlet predictions have been published. The most comprehensive reporting of the fit of the model is provided by Ehrenberg and Uncles (1999) who summarise the fit of the model for a range of predictions, for the leading eight brands in twelve product categories, across a variety of countries.

The main summary of empirical evidence from Ehrenberg and Uncles (1999) is reproduced below in Table 2.3.

Table 2.3
Annual Performance Measures for the Eight Leading Brands.
From Ehrenberg and Uncles 1999.
Observed [O] and Theoretical [T], for individual brands across a dozen product categories.

	Brand Share	% Buying		Purchases Per Buyer		% Buying Once		100%-Loyal % of Brand Purchases		Purchase Per Sole Buyer		Category Purchases per Buyer	
	O	O	T	O	T	O	T	O	T	O	T	O	T
Any Brand	100	(83)	**(83)**	(11)	**(11)**	12	**13**	100	**100**	11	**11**	11	**11**
First	(27)	46	**48**	4.6	**4.5**	36	**36**	20	**15**	4	**3**	13	**13**
Second	(19)	36	**37**	4.1	**4.0**	40	**41**	15	**12**	4	**2**	14	**14**
Third	(12)	26	**27**	3.7	**3.6**	47	**45**	10	**10**	4	**2**	15	**14**
Fourth	(9)	24	**23**	3.3	**3.4**	50	**47**	10	**9**	4	**2**	14	**15**
Fifth	(7)	16	**18**	3.7	**3.3**	46	**48**	11	**9**	3	**2**	15	**15**
Sixth	(5)	15	**14**	3.0	**3.3**	55	**50**	9	**8**	3	**2**	15	**15**
Seventh	(4)	13	**13**	3.2	**3.3**	55	**50**	8	**8**	3	**2**	15	**16**
Eighth	(3)	8	**9**	3.9	**3.2**	55	**52**	7	**8**	3	**2**	16	**16**
Average	(8)	(23)	**23**	3.7	**3.6**	48	**46**	11	**10**	3	**2**	15	**15**

( ) = Inputs to the model *Sources*: Nielsen, IRI, AGB, GfK, TCI

Ehrenberg and Uncles go on to summarise the fits as being "clearly close, with no average biases (except for the consistently low purchasing rates of the 100 percent-loyal buyers) and correlations between the observed and predicted values of about 0.9 or more. They conclude that the measures are therefore not idiosyncratic for specific brands. However, it is not clear how they can reach this conclusion with averaged data across many brands.

A number of other studies have also analysed the fit of the model, although they report less comprehensive results, usually for only a limited set of predictions.

Uncles et al. (1994) investigate the fit of Dirichlet predictions of average purchase frequency (w) and share of requirements (SoR) for 310 brands with market shares of more than 1 percent from 34 product categories in the US. The results of their investigation are reproduced in Table 2.4.

Table 2.4
Predictions of Purchase Frequency and Share of Requirements.
Adapted from Uncles et al. (1994).

Brands	Purchases per Buyer		Share of Requirements	
	Observed	**Theoretical**	Observed %	**Theoretical %**
average for 16 products with at least 10 itemised brands.				
1[a]	3.1	**2.9**	39	**38**
2	2.7	**2.6**	33	**34**
3	2.9	**2.5**	31	**33**
4	2.4	**2.3**	29	**32**
5	2.4	**2.3**	27	**31**
6	2.2	**2.3**	29	**30**
7	2.1	**2.2**	25	**30**
8	2.1	**2.2**	25	**30**
9	2.1	**2.2**	23	**29**
10	2.1	**2.2**	28	**29**
Average	2.4	**2.4**	29	**32**
Average r	.94		.93	
Average for products with < 10 itemised brands:				
Average	2.3	**2.2**	37	**43**
Average r	.71		.86	

In the 16 categories with more than 10 brands, the correlations for share of requirements and average purchase frequency are all high, averaging at 0.9. For the 18 categories with less than 10 itemised brands, correlations of about 0.7 for predicted average purchase frequency and 0.9 for share of requirements are reported. They say that the relatively low correlation for average purchase frequency is mainly because the values do not vary much.

Bhattacharaya (1997) reports that across 372 brands and 34 product categories a brand's actual share of category requirements (SCR) levels bear correlations of 0.77 with Dirichlet predicted SCR's. No information is given about the product categories over which this examination is conducted.

A different aspect of the fit of the Dirichlet model was investigated by East and Hammond (1996). East and Hammond examined the fit of repeat purchasing predictions over a number of analysis periods, instead of the conventional subsequent period. Across nine product categories East and Hammond reported an average over-prediction of repeat purchasing levels of

15 percent when following the behaviour of a cohort of buyers for a year. Or in other words 85 percent of repeat buyers stayed with a brand over a year.

Less systematic investigations of the fit of the model have also been reported, although these are usually only for individual product categories or for which the data are not presented. Morrison and Schmitlein (1988) note that the "model fits histograms of purchases extremely well, and that NBD forecasts for future purchasing patterns are often very accurate", pg 145. Barnard et al. (1994) report a correlation of 0.99 for the data of Lenk et al. (1993) for the proportions of buyers buying in a subsequent period conditional on them being non-buyers, once-only buyers and two-plus buyers in the previous period. For repeat buying proportions for consecutive purchases using Fader and Schmittlein's (1993) data, Barnard et al. also report a correlation of 0.7, while for share of category predictions they report a correlation of 0.997. Ehrenberg et al. (1990) report that the correlations with the observed values are about 0.7 or 0.8 across a wide range of measures in an illustrative application of the Dirichlet model to the instant coffee market.

Fader and Schmittlein (1993) examine the phenomenon of excess loyalty for some occasional high share brands, i.e., higher levels of repeat-buying than predicted by the Dirichlet. They also investigated share of category requirement predictions. They found that loyalty was under-predicted by the Dirichlet in most cases. However, as Barnard et al. (1994) point out, the average error was only 0.1, and also does not show whether this small average effect was due to one or two fairly large deviations for high share brands (or indeed deficits for low share brands). Barnard et al. (1994) go on to point out that the largest discrepancy was only 5 percent, for a predicted loyalty value of 72 percent. In summary, Barnard et al. (1994) conclude that the analysis of Fader and Schmittlein (1993) does not seem to imply many consistent and sizeable deviations from the Dirichlet predictions.

However, while these studies report correlations between observed and predicted quantities, there is little information about the dispersion of deviations between the observed and predicted values of the characteristics predicted.

Most of the studies support the hypothesis that the predictions are unbiased. The well known exceptions are the under-prediction of purchasing rates of the

100 percent-loyal buyers (Ehrenberg and Uncles 1999), and the over-prediction of repeat buying in long analysis periods (East and Hammond 1996).

Uncles et al. (1994) report that there are some isolated outliers (with as much as 50 percent difference between observed and predicted figures), and without these outliers the correlations between the observed and predicted values within their categories would return to 0.9 (sic).

One of the few studies to investigate the empirical distribution of deviation scores for predictions obtained from the Dirichlet model is Bhattacharaya (1997). Reporting on observed share of category requirements and Dirichlet predicted computed values Bhattacharaya concluded that the Dirichlet does a very good job, with 70 percent of the deviations in the ± 10 share of requirements points range.

Dekimpe et al. (1997) found that when studying the over-time behaviour of brand-loyalty, for 21 product categories, there is little support for the often heard contention that brand-loyalty continues to decline, and that while the variability around a brand's underlying loyalty level is not negligible it has not increased systematically over time.

Kahn, Kalwani and Morrison (1988) utilise the generalisation that the average amount bought per buyer times the proportion of non-buyers equals a constant across brands within a product category. While noting that this empirical generalisation is independent of the gamma-Poisson assumptions of the NBD model they concluded that in a simulated NBD world with typical values for consumer heterogeneity the variation of the constant across brands was not great, and thus is an appropriate benchmark to assess deviations from the norm. Studying 18 brands from four product categories they found evidence of 10 brands having deviations in average amount bought per buyer of more than 10 percent which they used to support their possible identification of niche or change of pace brands, although they note that there are other possible explanations.

According to Light (1998), the Ehrenberg tradition, which he labels the old view of marketing (specifically that measures of loyalty are correlated with market share) is insufficient. While conceding that Ehrenberg has replicated again and again that there is a correlation between market share and loyalty, Light maintains that brands deviate significantly from the Ehrenberg's "averages". As evidence for this assertion he reports results for two categories from Baldinger and Rubison's (1996) data, reproduced in Table 2.5. Light points out that brands B and C in Category One have markedly different shares, but similar share of requirements. In Category Two, Brands A and B have significantly different shares but similar share of requirements; brands B and C similar shares yet dissimilar share of requirements.

Table 2.5
Market Share and Share of Requirements.
From Light (1998).

	Share	Share of Requirements
Category One		
Brand A	43	63
Brand B	19	46
Brand C	7	46
Category Two		
Brand A	47	63
Brand B	18	62
Brand C	16	48
Brand D	11	42

The failure to observe a Double Jeopardy pattern (i.e., a correlation with market share) in the categories, however, could simply be because the categories are from partitioned markets.

The reporting of the average fits of the model as in Table 2.3 goes some way to establishing that the model provides largely unbiased estimates except for the noted bias in the prediction of buying rate for 100 percent loyal buyers. However, this information does not allow us to make any conclusions about the precision of the estimates for an individual brand. Applications of interpretive norms such as those employed by Sharp and Sharp (1997) in the investigation of a loyalty card scheme require that the model provide precise estimates.

Conceptually, large deviations would cast doubt on the managerial and scientific usefulness of the Dirichlet obtained predictions. However this criticism has to be balanced. Even in the situation where correlations between observed and predicted values are not particularly high, and considerable discrepancies between the observations and the predictions may occur, the choice is between competitive models, of which none seems to offer a reported similar order of magnitude of prediction, Ehrenberg (1993).

2.5.2 Known Systematic Discrepancies

Given the parsimonious input requirements of the model, it is not surprising that there are a number of systematic discrepancies. However, Ehrenberg et al. (1990) caution that when examining any apparent deviation, one should establish whether the discrepancy occurs consistently. They go on to note that typically, much work is usually needed just to diagnose and properly interpret a deviation.

However, a number of systematic deviations have been noted in addition to the bias conceded in the previous section for the buying rates of 100 percent loyals. Barnard et al. (1994) set out examples of systematic discrepancies such as:

(1) Over-predicting the number of customers who buy a brand more than once a week, for those brands which have significant number of once a week buyers (Ehrenberg 1959; Chatfield et al. 1966).

(2) Somewhat over-predicting the period to period repeat buying in relatively long analysis periods.

(3) Not allowing for the steady erosion of repeat buying observed over longish periods of time (East and Hammond 1994, Barnard et al. 1994).

(4) Under-predicting the average purchase rates of 100 percent loyal buyers, by up to a purchase or two (Ehrenberg 1988).

(5) Failing to provide an adequate fit to the observed distribution of category purchases, (Ehrenberg 1988, Uncles 1989).

(6) Occasionally under-predicting repeat buying for some high share brands (Ehrenberg 1988; Ehrenberg et al. 1990), and occasionally over-predicting repeat buying for some lower-share brands in a category (e.g., Uncles et al. 1994).

According to Ehrenberg and Uncles (1995), such deviations are due mainly to "other factors" which are excluded from the model specification. Furthermore, they add that the discrepancies are mostly not well documented, let alone understood.

2.6 Summary

The Dirichlet model is a stochastic approach to the modelling of consumer behaviour. By making probabilistic assumptions about purchase incidence and brand choice the model specifies the probability of a consumer making *n* purchases in a time period and which brand is bought for each of the *n* purchases. By aggregating these probabilities over consumers the model is capable of describing and predicting a number of aspects of aggregate consumer behaviour of interest to marketers.

A considerable number of markets have been analysed, and these markets are summarised by Ehrenberg as showing patterns consistent with the Dirichlet model.

On clarifying the fit of the model Ehrenberg and Bound (1993) explain:

> Fit to us depends on not merely establishing a significant (that is probably non-zero) correlation for a single data-set (SSoD). Instead, it turns on the absence of systematic bias across many sets of data (MSoD) and on correlations between the observed and theoretical values generally on the order of 0.9 , pg. 98.

CHAPTER THREE

ASSUMPTIONS OF THE Dirichlet model

3.1 INTRODUCTION

The development of Ehrenberg's tradition of modelling of aggregate patterns in the behaviour of consumers culminates in the comprehensive Dirichlet model for stationary and unpartitioned markets.

The Dirichlet model follows from the application of the negative binomial distribution to describe purchase incidence, first described by Ehrenberg (1959). In addition, Ehrenberg (1988) sets out a number of other general (and sometimes more easily applicable) approximations such as the Logarithmic Series approximations or the regularity, *average purchase frequency(1-penetration) = constant*, as set out in for example in Ehrenberg(1988). While earlier models give similar predictions to those of the Dirichlet model, the Dirichlet enables a greater range of measures to be predicted, mainly stemming from the incorporation of brand choice into the description of behaviour.

The Dirichlet model is a stochastic formulation of purchase incidence and brand choice at the individual level with very parsimonious requirements. The stochastic formulation of the model assumes a mixture of distributions at four levels:

(i) Purchasing of the product class takes the form of a Poisson process for each consumer,
(ii) The purchasing rates of different consumers are distributed according to a gamma distribution,
(iii) Each consumer's choices among the available brands follow a multinomial distribution,
(iv) These choice probabilities follow a multivariate Beta or "Dirichlet" distribution across consumers.

3.2 THE PURCHASE INCIDENCE MODEL

Table 3.1 depicts the purchase incidence component of the model. Briefly, it assumes that a consumer's purchase quantity in a given time period is a random variate generated from a Poisson process with mean rate μ. The individual consumer's mean purchasing rates (μ) are then assumed to be distributed across the population according to a gamma distribution. The combination of the Poisson and gamma distributions, gives rise to a negative binomial distribution of consumers making *1,2,3,..n* and so on purchases in a given time period.

Table 3.1
A stochastic representation of consumer purchases (x).
Adapted from Ehrenberg (1959).

Consumers	Period					Long-run Averages	Distribution
	I	II	III	IV	...		
A	x	x	x	x	x	μ_A	Poisson
B	x	x	x	x	x	μ_B	Poisson
C	x	x	x	x	x	μ_C	Poisson
D	x	x	x	x	x	μ_D	Poisson
...	x	x	x	x	x	$\mu_{...}$	Poisson
Mean	M	M	M	M	M	M	
Distributions	NBD for aggregate frequency of buying x					Gamma	

3.2.1 The Poisson Process

The Poisson assumption of the model implies that the number of purchases (x) of each consumer for a particular product by a single consumer (household) in successive time periods of equal length are independent and are drawn from a Poisson distribution with a long run mean μ. Thus, for a household which makes 5 purchases a year it is assumed that these purchases would be spread out randomly with a steady probability of buying of about 0.1 in any given week, and independently of precisely when the previous purchase was made. This assumption amounts to the specification of a Poisson distribution with a long run weekly buying rate μ of 0.1 for that household.

The stipulation of a Poisson process for purchase incidence means that a consumer's probability of purchase is unaffected by when they made a previous purchase, beyond that some consumers have different rates of purchasing (reflected by the gamma distribution of purchasing rates). This in turn implies that the individuals's inter-purchase times of a particular brand follow an exponential distribution. The resulting exponential distribution of interpurchase intervals has a mode at zero, which implies that likelihood of buying coffee today is unaffected by buying it in the previous period.

Thus the assumption clearly is not true in short time periods, where purchases made in one period would directly affect those made in the next. This is reconciled by Goodhardt et al. (1984) who state that the assumption of the Poisson process rests on the basic observation that purchase incidence tends to be effectively independent of the incidence of previous purchases, for periods greater than some minimum like a week.

Chatfield and Goodhardt (1973) state that a potential difficulty arises in the Poisson assumption when the product is not sold in clearly defined units. They give the example of the purchasing of petrol, where the units of purchase do not follow a Poisson distribution, however, they maintain that the NBD model can still be applied in such situations if the analysis is carried out in terms of purchasing occasions. Grahn (1969) also discussed such an application. There is no problem of analysing either purchase occasions or quantities, and in most cases the results are analogous, stemming from the observation that the average number of purchases on a given purchase occasion for most products is usually close to one, Ehrenberg (1988).

Review of the Poisson Assumption

The assumption of a Poisson process has been extensively tested, both directly, by for example, Ehrenberg (1959); Chatfield et al. (1966); Herniter (1971); Chatfield and Goodhardt (1975); Dunn, Reader and Wrigley (1983); Morrison and Schmittlein (1988) and indirectly throughout the generally successful reported applications of the NBD and Dirichlet models, see for example, Ehrenberg (1988). This justification needs to be considered further. As Ehrenberg (1995a) points out, whenever the Dirichlet predicts its wide range of aggregate measures, each measure is a hypothetico-deductive test of the model's underlying individual-level theoretical assumptions. However, Ehrenberg (1988) concedes that the fit, for example of the repeat-buying prediction, is not necessarily a very rigorous test of the underlying rationale and assumptions of a Poisson process with long run means distributed gamma across the population. This criticism obviously applies equally to the justification of other assumptions of the model by similar logic.

Initially, as a justification of the Poisson assumption, Ehrenberg (1959) examined the distribution of the differences of individual consumer's purchases in two equal time periods. Theoretically, these differences should be distributed with a standard deviation equal to $\sqrt{2m}$ where m is the mean amount bought by all consumers in each time-period, Chatfield et al. (1996). Chatfield et al. further add that this result is independent of the form of the compounding distribution.

Summarising for 25 brands with samples of about 2000 informants Ehrenberg (1959) concluded that the comparison of the theoretical ($\sqrt{2m}$) and the observed (root-mean-square) estimates, showed close agreement with no systematic deviations. Chatfield et al. (1966) add a further 70 comparisons and summarise "for standard deviations from about 0.1 to 1.5 the observed and theoretical values average at 0.55 and 0.59, and the mean absolute deviation of the individual differences is only about 0.1", pg 325.

A common criticism of the postulation of a Poisson process for consumer purchasing is identified, for example by Corlett (1966) who notes that he would expect more autocorrelation, either positive or negative than one would obtain from successive draws from a Poisson process.

Chatfield et al. answer this with the following.

> This not uncommon *a priori* doubt about the Poisson distribution is of some general importance because it must stem from a failure to understand Poisson distributions. They do not only describe rare events like Prussians being kicked to death. Instead, for large mean values they are better thought as describing a certain form of small irregular scatter about the mean. A *Poisson like* consumer with a mean of 100 will most likely buy between 90 and 110 units period after period; a consumer with a mean of 10 will buy mostly between 7 and 13 per

period; a consumer with a mean of 1 will buy 0,1, or 2 units per period and so on, (Chatfield et al. 1966, pg. 366).

The characterisation of individual purchase incidence as a Poisson process has also been criticised due to the derived exponential interpurchase times. The exponential distribution has its mode at zero, which means that the probability of purchase is the highest in the time period immediately following a purchase which seems intuitively unreasonable according to Murthi, Srinivasan and Tadikamallal (1993). However Ehrenberg (1988) points out that the often cited failures of the Poisson assumption (independence of purchases and a dead period following purchases), are characteristics of short term shopping and usage habits and have little if anything to do with people's longer term repeat buying and switching behaviour, and the basic finding is that the patterns tend to apply for periods of any length as long as they are relatively long compared with the minimum interpurchase time.

Replacing the Poisson Distribution

As identified above most of the criticism of the Poisson process results from the derived exponential distribution of interpurchase times. However, according to Morrison and Schmittlein (1988) assessing the regularity of an individual's interpurchase time is not easy. If the purchasing record is long enough to produce a statistically stable coefficient of variation the stationarity of the process may be in doubt.

The most often investigated alternative to exponentially distributed interpurchase times of a Poisson process, is an Erlang distribution (special case of a gamma distribution with an integer shape parameter). Erlang interpurchase times have a very simple relationship to the Poisson process. An Erlang-*s* random variate can be generated by recording every s^{th} arrival from an underlying Poisson process.

Chatfield and Goodhardt (1973) consider numerous comparisons between models assuming a Poisson process and an Erlang-2, the feature of interest being that the resulting distribution of interpurchase times has its mode at 1. The overriding conclusion from these comparisons is that the proposed modification does not produce histograms of purchases that are very different from a Poisson assumption. They go on to note that the purchase distributions of individuals with low average purchasing rates are virtually indistinguishable from Poisson distributions; only individuals with a high purchasing rate (about 6 units in 24 weeks) will have purchases which could be described by an Erlang distribution which would be clearly dissimilar from a Poisson distribution. They also note that such individuals only form a small part of the buying population. This does, however, highlight that in markets where the average rate of buying is high in comparison to the length of analysis, that the Poisson process assumption may be violated.

Dunn et al. (1983) also investigated the application of a Poisson process in the context of purchasing at individual stores and agree with the results of Chatfield and Goodhardt (1973), namely that the number of individuals for which the Poisson assumption can be rejected is small.

Chatfield and Goodhardt (1973) also consider repeat buying predictions for the gamma-Erlang-2 model as compared with the gamma-Poisson formulation and conclude that the differences between the two sets of predictions are small. They identify a possible explanation for this as being that the variance in the gamma distribution in the model is dominant, so that a failure in the model is therefore likely to be mainly due to a failure in the gamma assumption, thus the gamma-erlang model is unlikely to give a much better fit than the gamma-Poisson model, which they confirmed empirically.

In a similar study to Chatfield and Goodhardt (1973), Morrison and Schmittlein (1981) concluded that for four consumer non durables, forecasts using an erlang-2 model outperformed those of the NBD where the number of non-buyers was large. However, in their empirical comparison the conventionally formulated gamma-Poisson model outperformed the erlang-2 model in terms of conditional expectations for period to period repeat buying.

One of the most important aspects in interpurchase times which is not captured by the Poisson process was noted by Ehrenberg (1959), and also investigated by Dunn et al. (1983). The suggestion is that a substantial proportion of the population might exhibit weekly interpurchase times for frequently purchased package goods which would be inconsistent with draws from a Poisson process. Kahn and Morrison (1989) caution that the appropriateness of Poisson distribution for purchase incidence rests on whether weekly shopping patterns are observed or not. That is, if the distribution of interpurchase times is "spiked" for weekly purchases, then an erlang-2 specification may well be more appropriate, while the Poisson process fits the purchasing patterns of light buyers.

Wheat and Morrison (1990) take up this theme when examining individual household level data for coffee, and reported that across all households, coffee purchases are more regular than described by an exponential distribution, and thus are unlikely to be generated by a Poisson process. They go on to say that determining the regularity of consumers' purchases is an important - and often neglected - step in selecting the appropriate model of purchasing behaviour, and that if purchasing is more regular than the exponential assumption implies, that conditional expectations based on an erlang-2 gamma model may outperform those derived from a gamma-Poisson model. They however, do not test this empirically.

Despite these criticisms, Goodhardt et al. (1984) conclude that the evidence is that departures from Poisson purchasing are small and certainly not large enough to stop the overall model from working.

3.2.2 Gamma Heterogeneity

The gamma distribution has been traditionally used to account for the existence of different purchase rates across consumers. The gamma distribution is a flexible two parameter distribution that can assume a variety of shapes based on the values of the two parameters, Johnson and Kotz (1969).

In the Dirichlet model the long-run mean purchasing rates of different consumers (μ) are assumed to follow a gamma distribution (Γ), with a two-parameter density function:

$$\frac{e^{-\mu K/M}\mu^{K-1}}{(M/K)^K\ \Gamma(K)}$$

where μ is the long run mean purchasing rates for a consumer

M is the overall mean purchasing rate (per capita or per household)

K is the gamma distribution's second parameter and reflects the heterogeneity of consumer purchasing rates.

The results of aggregating the gamma-Poisson process across consumers yields a distribution of those buying 0,1,2,...x of the product which is distributed according to a negative binomial, Goodhardt et al. (1984) . The observation that the observed frequencies of consumers buying 0,1,2,...x fits a negative binomial distribution well has been used to justify the assumption of purchasing rates distributed according to a gamma distribution. However, this in common with the analogous justification of the Poisson process, is a weak test, since a number of different compound formulations can give rise to an aggregate negative binomial distribution of purchases. While conceptually there is no difficulty in seeing how well a sample of buying rates fit a gamma distribution, for example utilising a Komogorov-Smirnov test, this method of justification has not been used.

In addition, there is also a theoretical "characterisation" theorem for the gamma distribution (see Chatfield and Goodhardt 1973). This says in the present context that the distribution of the buying rates has inevitably to be a gamma if the following assumptions hold:

(1) That buying brand X is independent of buying brand Y
(2) That buying brand X is independent of X's share of the consumer's total product class purchases.

In practice the correlation between brands for assumption one tends to be in the order of 0.1, rather than the theoretically stipulated zero. And the second assumption is supported by the fact that market shares tend to be much the same among light, medium and heavy buyers (Ehrenberg 1988; Goodhardt et al. 1984; Shoemaker, Staelin, Kadane and Shoaf 1977). Thus, according to Ehrenberg (1988), since both assumptions 1 and 2 are approximately true in practice, then the stipulation of gamma distributed buying rates is approximately true. Ehrenberg goes on to claim that to model consumer heterogeneity in buying rates in this way is not an arbitrary assumption, but is both theoretically and empirically very well grounded.

The Zero Class Problem

The stipulation of gamma distributed rates of buying implies that as longer and longer time periods are considered eventually everyone will purchase i.e., all individuals have $\mu > 0$. However, clearly some consumers will never purchase a brand or product class. Frequently quoted examples where this assumption is violated are cigarettes and dog food, where never-buyers are non-smokers and non-dog owners respectively. This has special relevance when one considers that the estimate of the proportion not buying the product in a time period, but which never-the-less includes potential purchasers, is an integral quantity in the estimation of the model.

Morrison (1969a), for example, demonstrates that when genuine non-buyers exist this can lead to badly biased predictions for a subsequent period. Morrison also notes that for some product categories, there may not be a simple demographic characteristic that demarcates between users and non-users. Even where information about previous purchase histories is available, the issue of whether the absence of purchases represents a genuine non-user or a light user of the category is unresolved.

Morrison and Schmittlein (1988) reported that in the majority of applications of NBD derived conditional expectations that they have seen, the NBD under-predicts what the period one non-buyers will do in period two. However, they reported that two possible solutions (assuming a spike at zero, or assuming more regular purchases than Poisson) only exacerbated the problem. They go on to conclude that obviously some form of non-stationarity must be the cause of the under-predictions.

Lenk et al. (1993) repeat an earlier claim that the NBD model under-predicts the test-period purchases by previous non-buyers (Morrison 1969a; Morrison and Schmittlein 1981, 1988; Schmittlein and Morrison 1983). However, Barnard et al. (1994) point out that there seems to have been no published empirical evidence of such errors, and therefore any non-buyer discrepancy seems to be merely a theoretical possibility without any known empirical basis.

Complex heterogeneity

The assumption of gamma distributed purchase rates stipulates a smooth distribution of rates of buying for consumers. Dalal, et al. (1984) note that the heterogeneity of buying rates in the population of interest could in fact be quite complex, and in such cases, a simple gamma distribution may be inadequate. However, they also concede that most well known distributions are not capable of handling multiple modes and mass points.

One of the few studies investigating alternative purchase rate distributions is Murthi et al. (1993) . They investigated three alternative possible distributions for the rate of buying; inverse Gaussian, lognormal and Weibull distributions. They note, however, that all these distributions are fairly similar in shape when the coefficient of variation in the data is moderate (i.e, $\sigma/\mu \leq 1$). Their conclusion is that across low and medium levels of the coefficient of variation, the models perform equally well in fitting the data irrespective of the distribution, thus establishing the gamma distribution as a suitable approximation.

3.2.3 Negative Binomial Distribution of Purchases

The mixing of an individual level Poisson process with individual mean rate of buying μ distributed gamma across a population yields a negative binomial distribution (NBD) for the frequency of observing 0, 1, 2, 3, 4, ... and so on events. The NBD is always positively skewed, and for the values of parameters usually obtained from consumer purchasing data, the resulting distribution is reverse-J shaped, Ehrenberg (1959) and Grahn (1969).

The use of negative binomial distributions and models based on it have had a variety of previous applications, e.g., in the study of accident statistics, in certain ecological birth and death and contagious processes, and in some operational research theory, Chatfield et al. (1966). Applications to accident statistics in particular have been extensively reported, for example the classic paper by Greenwood and Yule (1920). In general, according to Ehrenberg (1988), the analysis of empirical data (in these typical applications) has been extremely limited, being concerned almost exclusively with the question of whether or not an observed frequency distribution can be fitted by a given theoretical distribution function.

The probability of n purchases follows a negative binomial distribution, with mean M, and exponent K:

$$p_n = (1+\frac{M}{K})^{-K} \frac{\Gamma(K+n)}{n!\Gamma(K)} (\frac{M}{M+K})^n$$

for n =1,2,3... ∞ purchases

Where

M =	average rate of buying in the population
K=	exponent of the negative binomial
Γ	denotes the Gamma Function,

The parameter K is calculated by fitting an NBD to the observed distribution of purchases for the whole product class. Techniques commonly used are summarised in § 4.3.1.

Ehrenberg (1959) states that the goodness of fit of the estimated distribution can in principle be assessed two ways, by calculating the value of χ^2 for the observed and theoretical frequencies, or alternatively a quicker test is to compare the observed variance with the theoretical value $m(1+a)$. Ehrenberg (1988) reports that a good fit to consumer purchasing data has been obtained in most cases, although these measures are seldom reported. Exceptions to this is Ehrenberg (1959), and the informal reporting of Chatfield et al. (1966). The latter paper reports that "the agreement between the observed standard deviations and theoretical variance of the fitted NBD was also generally good, for standard deviations up to 1 or 2", pg. 330.

As previously identified since a negative binomially distributed histogram of aggregate purchases for a population follows from compounding a Poisson process with a gamma distribution, the fit of the observed distribution of purchases to the negative binomial distribution may be considered as an indirect, but relatively weak test of the appropriateness of the assumptions. Dunn et al. (1983) comment that because alternative distributions are not considered the question answered is really only how well does the resultant NBD fit the observed data.

Discrepancies

A number of observations have been made about the suitability of the negative binomial distribution to describe the observed frequency of consumers buying 0,1,2... and so on times.

Chatfield et al. (1966) summarise that in general the NBD fits observed data well. However, for distributions of observed frequencies with relatively large standard deviations (about 5 or more), the theoretically predicted distributions standard deviations tend to be somewhat larger. This means that the observed distributions do not appear to quite as skew as a NBD. It has also been noted that distributions with large standard deviations tend to be those products which are heavily bought, as for example margarine, detergents and bread, Ehrenberg (1959).

There is, also a tendency for the observed proportions buying *n* times to be under-predicted for quantities near the number of weeks in the analysis unit, this is manifestation of regular shopping patterns, accounted for by a the Poisson assumption. Thus for products which tend to be bought at most once a week, the number buying more than 26 times in a 26 week period is more than the theoretically estimated value. This phenomena has been reported, for example, by Ehrenberg (1959), Chatfield et al. (1966) and Chatfield and Goodhardt (1973).

Alternatives

While the assumption of a negative binomial distribution to describe aggregate purchase levels has been questioned, few studies have attempted to investigate the performance of alternative distributions. Tjhis may be because of the mathematically tractable results obtained from the compound gamma-Poisson model.

One exception to this is Murthi et al. (1993) who consider three alternative candidates to the conventional negative binomial model.

(1) Condensed negative binomial (where the interpurchase time follows an erlang-2 distribution and heterogeneity is specified as gamma).

(2) Generalised negative binomial.

(3) Generalised Poisson .

Table 3.2
Purchase Incidence models and their properties.
Adapted from Murthi et al. (1993).

NBD Model	Parameters		
NBD	(a,k)	mean variance P_0	$= m = a\ k$ $= m(1+a)$ $= (1+m/k)^{-k}$
CNBD	(a, k)	mean variance P_0	$= m = a\ k$ $= m(1+a)$ $= m\ a+m/2+[1-(1+4a)^{-m/a}]/8$
GPD	(θ, λ)	mean variance P_0	$= \theta\ (1-\lambda)^{-1}$ $= \theta\ (1-\lambda)^{-3}$ $= Exp\ (-\theta)$
GNBD	(α, β, η)	mean variance P_0	$= m = \eta\alpha/(1-\beta\alpha)$ $= \eta\alpha(1-\alpha)\ /\ (1-\beta\alpha)^3$ $= (1-\alpha)^\eta$

P_0 = Proportion of the observations with zero purchases

Their conclusion is that across low and medium levels of the coefficient of variation, all models (except the CNBD) perform equally well in fitting the data irrespective of the heterogeneity distribution, and they conclude that overall there is no compelling reason to recommend any of the proposed modifications over the traditional negative binomial distribution. They do however, note a number of situations where alternative models fit better (for example when purchase incidence is Erlang-2 distributed then CNBD model performed better). While this addresses the theoretical possibility of using other distributions to describe purchases, most of the evidence presented elsewhere can be summarised as lacking the impetus to suggest the assumptions underlying the purchase incidence component of the model be rejected.

3.3 THE BRAND CHOICE MODEL

Following the description of the purchase incidence component of the model as set out in §3.2, this section sets out to describe how purchases are allocated across brands in the market.

The brand choice model assumes each consumer's choices among the available brands follows a multinomial distribution, and, that individuals' choice probabilities follow a multivariate beta or "Dirichlet" distribution across consumers. These assumptions are stated within the model more precisely as:

(i) An individual's brand choices over a succession of purchases are as if random, with a probability *(p)* of choosing brand *j*. These probabilities are fixed over time and brand choices at successive purchases are assumed independent.

(ii) The probability (*p*) of buying a given brand (*j*) vary amongst the individuals according to a Dirichlet distribution.

The integration of these assumptions concerning brand choice are depicted in Table 3.2.

Table 3.3
A stochastic representation of brand choice.

Consumers	Probabilities of buying brand					
	a	b	c	...	j	...
A	p_A	p_A	p_A	p_A	p_A	Multinomial
B	p_B					
C	p_C					
D	p_D					
...	$p_{...}$					
Mean	m	m	m	m	m	
Distributions (vertically)			Dirichlet			

3.3.1 Multinomial Process

According to Bass et al. (1984), although many are persuaded on the basis of theory or intuition that the brand choice process is characterised by purchase event feedback, there are a mixture of theoretical and empirical reasons to believe that consumer brand choice could be closely resembled by a zero-order process, and thus represented as independent multinomial probabilities fixed over time. Morrison and Schmittlein (1988) also conclude that the zero-order specification has substantial empirical support.

The brand choice aspects of the model that underlies the Dirichlet model is that in a stationary market a consumer has certain steady probabilities p_a, p_b, p_c, ... and so forth of buying brands *a, b, c* and so on. It assumes that the probability of purchase of brand *x* is constant over time and independent of prior purchase events. That is, an individual consumer habitually buys from a small repertoire of, for example, three brands with steady propensities or probabilities of say .6, 3, and .1 and with the other brands at virtually zero for that consumer. As an example, Hauser and Werenerfeldt (1990) report that in most grocery markets, brand portfolios (the number of brands a consumer has long term non-zero probabilities of buying) tends to be 4 or less, while Stern (1997) notes that in grocery markets the favourite brand typically accounts for about 65 percent of category requirements.

Zero-Order Process

In the Dirihclet model consumer's propensities are expressed as steady multinomial probabilities that operate independently over time, thus brand choice is assumed to be zero-order, for which there is much support, [according to, for example, Bass et al. (1976); Bass et al. (1984); Kahn et al. (1988); Ehrenberg (1988); Ehrenberg et al. (1994)]. In a two brand product category these probabilities reduce to a Bernoulli process, Morrison and Schmittlein (1988).

According to Ehrenberg (1995a), weight is given to the justification of a zero-order process through the close predictions of all the various aggregate measures of buying behaviour that have been tested. However, both Jones (1979) and Blattberg (1981) much earlier noted the need to investigate this fundamental precept of stochastic model building. Blattberg, for example, states that "the question of the order of purchasing behaviour to determine whether there is purchase feedback has not been answered satisfactorily", pg. 204.

Multinomial, Markov and Linear Learning models have all been proposed to represent brand choice behaviour. In the previously discussed multinomial approach an individual consumers' probability of purchasing brand *x* remains constant over time. The Markov model allows previous purchases to affect the consumers' probability of purchasing brand *x*. However, this is usually limited to a fixed number of previous purchases, thus distinguishing this approach from a Linear Learning model which allows all past purchases to affect the current probability of purchasing brand *x*.

Givon and Horsky's (1978) findings are representative of the numerous studies into brand choice. They found that individuals may differ in the order of the stochastic process they follow. Direct evidence of the process order has been provided through individual-type analyses of long purchase strings by Frank (1962), Massy (1966), Jones (1973), and Blattberg and Sen (1976). Givon and Horsky (1978) went on to investigate the application of different order processes and concluded that there was support for Linear Learning, Zero-order and Markov models as representations of brand choice.

Bass et al. (1984) summarised the literature concerning brand choice up until their study, and Table 3.4 presents the results.

Table 3.4
Summary studies on the order of the brand choice process.
From Bass et al. (1984).

Reference	Products tested	Sample Size	Test Unit	Test Used	Results and Comment
Kuehn (1958)	Snow Crop Orange Juice	13,519	Market	No Formal Statistical Test	Rejected zero-order hypothesis. Did not provide for heterogeneity in purchase probabilities.
Carman (1966)	Crest Toothpaste	5,028	Market	Constrained Regression Analysis	The Linear Learning model fit the data well.
McConnell (1968)	Experimental Beer Data	1,200	Market	X^2-test	The linear-learning model provides a good fit for only one of the three cases tested.

Massy, Montgomery & Morrison (1970)	Regular coffee, Kuehn's data, Folger's coffee, McConnell's data	531 13,519 805 1,200	Market	X^2-test on conditional expectations	Zero-order hypothesis rejected in favour of higher models for regular coffee and Kuehn's data. Results inconclusive for other products
Frank (1962)	Hill Brothers' Chase & Sandborn Coffee	71 53	Individual Individual	Runs Test Runs Test	Zero-order could not be rejected for approximately 75 percent of the families at 95 percent confidence level.
Massy (1966)	Regular Coffee	39	Individual	X^2-test	Zero-order could not be rejected for 85 percent of the families. Test carried out only on stationary families.
Wierenga (1974)	Food Product Beer Margarine	672 627 1,059	Individual	Runs Test	Zero-order rejected for a majority of the households
Blattberg and Sen (1976)	Facial Tissue Waxed paper Aluminium Foil	150 43 49	Individual	Bayes Approach	Found evidence for both zero-order and Markovian behaviour. Note that the models are not compared directly
Bemmaor (1979)	Regular coffee Instant coffee Margarine	677 531 685	Individual	Runs Test	Zero-order hypothesis could not be rejected in majority of the cases.
Givon and Horsky (1979)	Toothpaste Facial Tissue Waxed Paper Liquid Detergent Headache remedy	7,500 8009 873 751 771	Market	X^2-test	Found evidenced for both zero-order and learning.
Jeuland, Bass & Wright (1980)	Cooking Oil	1,683	Individual	Runs test	Zero-order hypothesis could not be rejected in 41 percent of the cases.

Despite the extensive investigation of this problem, there is still substantial disagreement about the issue of the order of brand choice. Unsurprisingly, the general conclusion is that it is possible that both zero-order and higher-order processes could be explain brand choice behaviour.

However, Massy (1966) identifies that the central issue facing modellers is whether departures from a zero-order process are consistently serious enough to warrant using a higher order model to describe brand switching behaviour. Dalal et al. (1984) also support Massy's call, and they go on to say that it is difficult to test the assumption about the order of the brand choice process in isolation. As a solution they advocate the use of a pragmatic approach, namely to recognize that all models are approximations to reality, and the question is whether they are adequate representations.

In addition, Bass et al. (1984) suggest that a possible explanation for the variation in results in studies of the order of the brand choice process is due to the level of aggregation employed in the analysis. They suggest that the heterogeneity in consumer brand purchase probabilities and a presence of non-stationarity (especially in purchase sequences that extend over a year or more) have confounded the assessment of the order of the brand choice process. According to them both heterogeneity and non-stationarity tend to produce results which suggest the switching process governing family brand choices is of a higher order than really is the case. They note that the non-stationary segment presents special problems for modelling, and thus believe that testing at the individual level, with appropriate regard to stationarity is the better method of examining the nature of the consumer choice process.

Bass et al. (1984) investigated nine product categories from a consumer diary panel (White Bread, Regular Coffee, Selected Soft Drinks, Margarine, Salad Dressing, Ready to Eat Cereal, Sugar, Ice Cream, Sausage) and a laboratory brand choice situation for Soft drinks. Their initial step was to classify households as either stationary or non-stationary on the basis of a chi-square test of the null hypothesis H_0: $p_{ij}^{m} = p_{ij}$ for all m (sequences of equal length, in this case they limited their investigation to the consideration of two equal lengths), where p_{ij} is the probability of switching from brand i to brand j over the entire purchase sequence.

Utilising four different tests (t, likelihood ratio, binomial runs and multinomial runs) they attempted to investigate the order of the choice process for the previously classified stationary households. They found that the majority of stationary households have purchase sequences that are consistent with a zero-order process or something close to it. In their data, about a quarter of the households were rejected as being non-stationary, however, even in these households a substantial proportion could still be characterised as being zero-order. They go on to conclude that there may be higher order processes in a population, but that the individual choice behaviour of a large number of consumers is consistent with the zero-order process, thus explaining why descriptions of aggregate behaviour which derive from zero-order assumptions are fairly accurate. In the context of the successful applications of Dirichlet modelling this suggests that the assumption of a zero-order process has not been excessively violated.

Schmittlein, Cooper and Morrison, (1993) suggest that some segments of the market might be better described by higher-order brand choice models. As one investigation of this idea, East and Hammond (1996) however concluded that the hypothesis that choice probabilities of light buyers changing more (and thus not being zero-order multinomials) could not be supported.

3.3.2 Dirichlet Heterogeneity of Choice Probabilities

The earliest attempts at modelling stochastic brand choice assumed that consumers were homogenous and that a single set of parameters could represent the buying behaviour of all consumers, Blattberg and Sen (1976). However, this assumption was soon rejected, since consumers were easily observed as having different rates of buying and different choice probabilities.

Within the Dirichlet model, this theme is taken up with consumers differing from each other in the size and composition of their brand repertoires and in their brand choice probabilities within these repetioires. The probabilities of buying various brands vary from consumer to consumer according to a particular form of multi-variate beta mixing distribution (known as the Dirichlet). The probability of choosing a particular brand reduces to the simple Beta-distribution when only considering two brands, Goodhardt et al. (1984).

The density function for the distribution of consumer's probabilities for choosing brand j is given by:

$$\Gamma(S)\ /\ \Pi(\Gamma\alpha_j)\ p_j^{\alpha_j-1}(1-p_j)^{(S-\alpha_j-1)}$$

where $p_j \geq 0,\ \Sigma p_j=1,\ S=\Sigma\alpha_j,\ and\ \alpha_j>0$

Techniques for estimating the parameter S are reviewed in section § 4.4.

Justification for the assumption of Dirichlet distributed brand probabilities is offered by Goodhardt et al. (1984). A requirement of the Dirichlet model is that a market must be unsegmented, which is to say that the proportion of purchases devoted to any particular brand is independent of the way the remaining purchases are distributed between the other brands. The crucial point is that the independence of the purchase probabilities, except for the constraint that the purchase probabilities for all brands must sum to one, is a characterisation of the Dirichlet distribution (Mosimann 1962, 1984).

Complex Brand Choice Heterogeneity

It is commonly advocated that consumers fall into distinct segments of loyals and various kind of switchers, rather than simply having varying degrees of split but habitual loyalty, as might be assumed by postulating Dirichlet heterogeneity of brand choice.

For example, Bass et al. (1976) note that the heterogeneity of the population as far as the choice process is concerned is much more complex than that for which any heterogenous multinomial model can account. Jones (1973) also proposed a model assuming complex heterogeneity in consumers' choice processes, where consumers could be represented by Zero-order, Markov and Linear Learning models.

However the justification for considering alternative heterogeneity formulations seems to be based on theoretical possibilities rather than empirical failure of the model. For example, Blattberg and Sen (1976) proposed that aggregate markets might be made up of a number of different segments each with different models, and requiring different parameterisation for heterogeneity. As justification for such an approach they note:

> Though such models (assuming the same heterogeneity distribution across the population), e.g., [3,14], can be useful in predicting repeat purchasing and brand switching at the *aggregate* level, they typically do not provide an accurate picture of the degree of heterogeneity that may be present in the market, (Blattberg and Sen 1976, pg 43).

Following a cluster analysis procedure they showed that there could be nine different segments each requiring different stochastic brand choice models. However they fail to develop why we should be interested in heterogeneity as an independent characterisation, or why, given the successful predictions obtained from simpler models, that such modifications are to be preferred, except for their emphasis on the possibility that one might want to make disaggregate level predictions.

Despite this, applications of the assumption of distinct segments have been incorporated into theoretical models of buyer behaviour (see for example, Massy et al. 1970; Bass 1974; Blattberg and Sen 1976; Colombo and Morrison 1989; Kamkura and Russell 1989). These models require specific parameter estimates for each segment. However, according to Ehrenberg (1995a) no generalisable results seem to have been reported which might give compelling reasons for such developments.

Barnard et al. (1994) highlight one of the difficulties in identifying appropriate heterogeneity distributions with the example of the 100 percent loyal customers. They point out that if consumers are polygamous, some will happen to appear monogamous in some chosen analysis period, and that this is closely predictable. They conclude that there is no systematic evidence in the literature that homogenous, distinct and substantial segments of "loyals" or "switchers" can clearly be identified.

3.4 Summary

The Dirichlet consists of a set of assumptions about purchase incidence and brand choice. These assumptions are represented by the use of the Poisson process to describe purchase incidence for an individual consumer, and the long run rates of purchasing incidence is distributed gamma across consumers. Brand choice is represented by multinomial probabilities for different brands in the market, and consumers differ from each other in the size and composition of their brand repertoires and in their brand choice probabilities within these repertoires according to a Dirichlet distribution.

While a number of alternative representations of purchase incidence and brand choice have been suggested in the literature, only a limited number have been investigated systematically. None seem to have the empirical support presented by Ehrenberg and his colleagues. Thus while a limited number of specific tests of the assumptions underying the model exist, a mixture of empirical and theoretical reasons are said to justify the assumptions. Primary amongst the reasons is the successful integration of a number of empirical observations across a wide range of product categories.

CHAPTER FOUR

ESTIMATION OF PARAMETERS OF THE NBD/Dirichlet model

4.1 INTRODUCTION

The formulation of the Dirichlet model as a set of stochastic assumptions about consumer purchasing as set out in Chapter Three offers a number of ways of obtaining predictions from the model.

For example, one approach might to be to simulate a panel of consumer purchases. In this approach, simulated consumer purchases would be generated based on the Dirichlet distributions with parameters obtained from the observed data. The distribution of purchases, the proportion only buying a single brand, the buying of the category and so on could then be tabulated.

The conventional method is to use the negative binomial distributed aggregate purchasing of the product category, and then allocate purchases across the brands. Predictions from the model then require the summation of the probabilities of purchasing a particular brand, conditional on having made a certain number of purchases in the product class.

Following the notation of Goodhardt et al. (1984), the main simplification is to reduce the calculations to that of a Beta-Binomial distribution. The proportion of consumers buying the product class n times and buying brand j, r times is given by the product of:

$$p(r_j, n) = P_n \, p(r_j|n)$$

where P_n = probability of buying the product class n times

$p(r_j, n)$ = probability for those of buying brand j r times out of n

P_n is given by

$$P_n = (1+\frac{M}{K})^{-K} \frac{\Gamma(K+n)}{n! \; \Gamma(K)} (\frac{M}{M+K})^n$$

for $n = 1,2,3 \ldots \infty$ purchases

where M is the average rate of buying in the population.
K is the heterogeneity of rate of buying in the population.
Γ denotes the Gamma Function.

The parameter K is calculated by fitting an negative binomial distribution of purchases for the whole product class. Techniques available for this are summarised in § 4.3.1.

The second term is given by:

$$p(r_j|n) = \binom{n}{r_j} \frac{B(\alpha_j + r_j, S - \alpha_j + n - r_j)}{B(\alpha_j, S - \alpha_j)}$$

where B denotes the Beta Function.
j denotes the parameter for the j[th] brand.
n is the number of purchases of the product class, and
r is purchases of a brand, $r \leq n$.
S the brand choice heterogeneity factor.

and $\alpha_j = S \frac{m_j}{M}$, where $\frac{m_j}{M}$ is the market share of the j[th] brand,

and $S = \sum_{j=1}^{j} \alpha_j$

Suggested solutions to the problem of estimating the parameter S are summarised in §4.4.1

4.2 BACKGROUND

In the context of the Dirichlet model, a number of authors have noted that there are many matters of estimation and goodness of fit which need to be considered, e.g., Bartholomew (1984). However, this is at odds with the view noted by for example, Goodhardt et al. (1984), who claim that the main justification of the Dirichlet model is that in practice it fits many different aspects of buying behaviour under a wide range of conditions.

Bemmaor (1995) also claims that the quality of interpretation depends on the accuracy of the parameter estimates. According to Bemmaor "more work needs to be carried out on the estimation issue to enhance the understanding of data", pg 114.

Reflecting on the choices available in the estimation of the model, Morrison and Schmittlein have repeatedly identified the need for the development of statistical properties for predictions, as a function of the estimation method selected, (see for example, Morrison and Schmittlein 1981; Schmittlein and Morrison 1983).

Morrison and Schmittlein (1981) go on to state "without some estimate of sampling error no rejection of the NBD as a predictive models has been attempted,. . . without a comparison of the efficiency of those estimates any inferences are left open to question", pg 1021.

4.3 PRODUCT CLASS PURCHASING

4.3.1 ESTIMATION

The theoretical NBD distribution of product class purchases in a period can be generated by:

$$P_n = \left(1+\frac{M}{K}\right)^{-K} \frac{\Gamma(K+n)}{n!\ \Gamma(K)} \left(\frac{M}{M+K}\right)^n$$

for n = 1,2,3... ∞ purchases

where	M	is the average rate of buying in the population.
	K	is the heterogeneity of rate of buying in the population.
	Γ	denotes the Gamma Function.

Morrison and Schmittlein (1988) summarise two common estimation methods. The first, often called the method of moments (MoM), uses the mean and variance of the observed of the distribution of the number of purchases. The second estimation method usually called the method of means and zeros (M&Z), uses the mean and the observed proportion of non-buyers from the observed distribution of the number of purchases

In both methods, the estimate of the mean is obtained from the sample mean, that is the average rate of buying of the product class.

Despite maximum likelihood estimates for the NBD being given implicitly by Anscombe (1950), few studies have utilised this method of estimation. One possible reason for this is that the maximum likelihood equations for K are very cumbersome to solve, Ehrenberg (1988). Morrison and Schmittlein (1988) also discount the use of this estimator since it involves a numerical line search.

Method of Means and Zeroes (M&Z)

Goodhardt et al. (1984) note that if the distribution of purchases for the whole product class is reverse J-shaped, and fitting by the method of mean and zeroes has high efficiency (e.g., Anscombe 1950; Ehrenberg 1959, 1988; Chatfield 1966). If the distribution of purchases is not reverse J-shaped, then the method of moments estimators could be used.

The means and zeroes method attempts to minimize the difference between the observed and predicted value of proportion not buying (P_0), (hence the "Zeroes" appellation), by solving for K.

Thus the proportion not buying (P_0) is:

$$\min_K \; P_0 - \hat{P}_0$$

Which can be written:

$$\min_K \; P_0 - \left(1+\frac{M}{K}\right)^{-K}$$

where M is the average rate of buying in the population.
K is the heterogeneity of rate of buying in the population.

Method of Moments Method (MoM)

The method of moments estimator requires equating the observed sample variance of the number of purchases (n) to its expected value:

$$M\ (1+\frac{M}{K})$$

by solving for K,

where M is the average rate of buying in the population.
K is the heterogeneity of rate of buying in the population.

According to Ehrenberg (1988) this method of estimation is not particularly efficient. Earlier, Ehrenberg (1959) claims that the efficiency is sometimes less than 50 percent. Ehrneberg (1988) claims that it is "labourious to compute the sample variance especially since in market research the basic frequency distribution is often not tabulated," although this seems less relevant given modern computing technology.

Theoretical Distribution of Purchases

Once estimates of the M and K are obtained, the theoretical distribution of purchases can then be generated by the formula for P_n .

$$P_n = (1+\frac{M}{K})^{-K}\ \frac{\Gamma(K+n)}{n!\ \Gamma(K)}\ (\frac{M}{M+K})^n$$

for $n = 1,2,3... \infty$ purchases

where M is the average rate of buying in the population.
K is the heterogeneity of rate of buying in the population.
Γ denotes the Gamma Function.

However since the possible number of purchases that a consumer makes is potentially unbounded, it is impossible to calculate all possible values of P_n. The usual procedure is to use an approximation procedure to sum the tail of probabilities over two values of n, (n' and $n'+1$), Ehrenberg (1988), pg 338. The position of the tail and the probabilities attached to these values of n are found by solving the equations:

$$P_{n'} + P_{n'+1} = 1 - \sum_{n=0}^{n*} P_n = P_R$$

$$n'P_{n'} + (n'+1)P_{n'+1} = M - \sum_{n=1}^{n*} nPn = Q_R$$

$$n' = integer\ (Q_R/P_R)$$

At this stage, $P_{n'}$ and $P_{n'+1}$ are able to be calculated.

$$P_{n'+1} = P_R\ (\ \frac{P_R}{Q_R} - n'\ )$$

$$P_{n'} = P_R - P_{n'+1}$$

Thus the cumulative number of purchases $\sum_{n=1}^{n'+1} n\ P_n = M$

Ehrenberg (1988) points out that it is possible to calculate a χ^2 statistic at this stage to assess the goodness of fit between the observed and predicted distributions of purchases. Expanding on this, where the observed product class distribution shows deviations from a NBD distribution, it is possible to substitute the observed product class distribution of purchases for the theoretically derived distribution. This case gives rise to the empirical-Dirichlet model, Goodhardt et al. (1984).

4.3.2 COMPARISONS

Relatively few studies have compared the effect of different estimation techniques on either the fit of the distribution of purchases or the subsequently derived statistics.

Morrison and Schmittlein (1981) note that the method of means and zeroes relies heavily on the observed proportion of zeros and that the difficulty in estimating this proportion can be expected to have a greater effect on predictions with those estimates arrived at using the method of moments. Morrison and Schmittlein went on to make comparative predictions using mean and zero estimators and method of moment estimators for four consumer non durables. They concluded that the predictions using the mean and zeros estimates are never the best. Finally the method of moments show the smallest deviations in two cases for the NBD, with the method of moments estimators for an Erlang-2 model working better when the number of non-buyers is large.

Dalal et al. (1984) noted that estimates obtained for the NBD model using MoM and M&Z methods were found to be almost the same. The apparent difference between the findings of Morrison and Schmittlein and those of Dalal et al. (1984) can be reconciled when one considers that Morrison and Schmittlein were concerned with the theoretical possibility of a large proportion of never-buyers whereas Dalal et al. were concerned with estimates obtained from real data.

Murthi et al. (1993) investigated the method of fitting in a simulation study. Across seven different values of the coefficient of variation for the distribution of average rate of buying, the method of means and zeroes always performed better (in terms of chi-square values). They also noted that when the coefficient of variation was low, the two methods of estimating model parameters yielded similar results.

4.4 BRAND CHOICE PROBABILITIES

4.4.1 ESTIMATION

Goodhardt et al. (1984) set out a specific procedure to estimate the parameter *S*, although it must be noted that other methods exist. The method detailed here is the method suggested by Ehrenberg (1988) and used in the commonly available BUYER software, Uncles (1989).

The parameter *S* is used to describe consumer's brand choice process. Formally it describes the the proportion of consumers making *r* purchases of brand *j*, conditional on them making *n* purchases of the product class. As previously identified this is an intermediary step, since to obtain predictions from the model the conditional probabilities need to be summated since there is no explicit formulae linking the various theoretical statistics required and *S*.

According to Goodhardt et al. (1984), the simplest estimating statistic to use appears to be the number of non-buyers of the brand. By the additivity property of the Dirichlet, the distribution of purchases of brand *j* reduces to a Beta-Binomial distribution for "brand j" or "not brand j". For consumers buying the product-class n times in the analysis period, the conditional probability of *not* buying brand *j* is therefore

$$p(0|n) = \frac{(S-\alpha_j)\ (S-\alpha_j+1)\ \dots\ (S-\alpha_j+n-1)}{S(S+1)\ \dots\ (S+n-1)},\ for\ n \geq 1$$

where B denotes the Beta Function.
j denotes the parameter for the j^{th} brand.
n is the number of purchases of the product class, and
S the brand choice heterogeneity factor.

and $\alpha_j = S\dfrac{m_j}{M}$, where $\dfrac{m_j}{M}$ is the market share of the j^{th} brand,

and $S = \sum_{j=1}^{j} \alpha_j$

Multiplying this by the probability P_n of making n purchases of the product class gives the overall probability of not buying brand *j*. This can then be equated to the observed proportion of non-buyers of brand *j*, and solved for *S*.

This obtains separate estimates of *S* for each brand. Goodhardt et al. (1984) adopt the convention that estimate for the product category is a weighted sum, with the weights according to market share.

4.4.2 ALTERNATIVES AND COMPARISONS

Goodhardt et al. (1984), note that the computational procedure set out above is "certainly is not the last word", (pg 654). For example, Goodhardt et al. (1984) identify that one particular aspect of the estimation of *S* that needs further study is the constancy of *S* over different time periods. They note that the parameter *S* in the model should be constant irrespective of the length of time period considered, but it varied from 1.0 when estimated for 4 week data to 2.2 for the 12 week data, to 1.8 for the year, in a study of the buying of toothpaste brands. Not withstanding this, they note that the predictions generally work well.

A number of authors have identified alternative methods of estimating the parameter S. With reference to Goodhardt et al. (1984), Kemp (1984) and Zufryden (1984) discuss alternative methods of estimation, while in a reply to Ehrenberg (1995), Bemmaor (1995) provides details of another alternative method.

Kemp (1984) states that he is unimpressed by the Goodhardt et al. (1984) estimation procedure for S. Instead he advocates the use of a method of moments estimator for the marginal distributions, rather than an estimator that uses just the marginal probability for not buying a given brand.

Zufryden (1984) criticises the estimation procedure as having computational and estimation difficulties. Expanding on this, Zufryden comments that the estimation procedure involves an iterative method that attempts to estimate the parameter S independently for the proportion not buying each brand and then reconciles any differences by, rather arbitrarily, weighting the derived S estimates to obtain the unique S parameter for the category.

Zufyden (1984) goes on to suggest an alternative estimation method:

> First estimating the NBD parameters, inserting these and replacing α_j by $S\ m_j/M$ in the P_0 equation of each brand and then minimizing the squared deviations between theoretical and observed P_0, summed over each brand j, by choice of S. This least squares method would still require truncation of the P_n terms and iterative search for a globally optimal S value. The individual α_j 's could then be obtained from the brand market share equations, pg. 651

Discussing the methods suggested by Kemp and Zufryden , Goodhardt et al. (1984) claim that one advantage of their method is that the observed and predicted brand shares are constrained to be equal, something which makes the model more attractive to potential users. They go on to state that this requirement is not met by Zufyden's method. Furthermore, they claim that they see no reason to minimise the sum of squared deviations between observed and predicted proportions not buying a brand as per Kemp's method, the reason given being that this proportion varies markedly.

Goodhardt et al. (1984) also discuss Kemp's suggestion of using the first two moments of the marginal distributions, but dismiss it because the observed second moment is not usually reported in marketing data used in the analysis, although they note that in principle it would be possible to calculate from raw data. However, they note that as far as the fitting of the NBD model is concerned for reverse J shaped distributions, it is better to fit by mean and zeroes, than by moments.

Bemmaor (1995) claims that a simultaneous search for a complete set of parameters leads to improvements in the model's predictions. He advocates a search for S, such that the differences between the observed and predicted proportions buying 1, 2, 3 ... of each brand are minimized. Developing this suggestion, Bemmaor gives as an illustrative example, the average duplication of Shell with other oil companies obtained by his estimator, 71 percent versus 72 percent observed; whereas Uncles and Ehrenberg (1990a) predicted 100 percent).

However Bemmaor (1995) fails to consider the fit of the model across other brands, which reproduced below demonstrates that while Bemmaor's estimate improves the fit for Shell, it is worse for the remaining brands.

Table 4.1
'Constant D' approximation predictions of duplication and Bemmaor's estimator.

	Shell	'others'	BP	Total	Mobil	Esso	Chevron
Observed	72	74	61	48	48	47	33
Uncles & Ehrenberg	100	72	62	39	40	40	27
Bemmaor	72	51	44	28	28	28	19

Hariharan and Velu (1993) note that only a few researchers and only recently have investigated the estimation of the parameters of the Dirichlet distribution. They identify three methods as being proposed in the literature for estimating the parameters of a Dirichlet distribution; a maximum likelihood (ML) approach, an approximate method based on geometric means, and a method of moments approach. After discounting the use of the method of moments method, due to the results of Narayanan (1991) who concluded, that even in small samples, the ML approach performed better in terms of bias ratio and squared errors. They found that the extent of the basis in the ML estimator increases as the sample size decreases, and that the final parameters are always biased upwards. They also found that when two the parameters have the same value the bias virtually explodes.

4.5 SUMMARY

In a reply to Bemmaor (1995), Ehrenberg (1995) notes that the model's predictions of such things as penetration, purchase frequency, sole buyers, and duplication of purchase and so on, should be the focus of any evaluation of different estimators. The question needs to be asked whether any estimate of *S* works better in all respects and always, if not, he suggests there is little point in worrying about different estimates of *S*.

Nevertheless, Ehrenberg (1994) concedes the importance of *S*, for example, he noted

> that one needs to understand this variation and the meaning of S more generally. How far is the variation in *S* mere sampling? What other explanatory factors are involved? What effects does the variation in S estimates have on the predictions of the model? pg 130.

Barnard et al. (1994) also note that an important outstanding issue is the variability and detailed meaning of the Dirichlet parameter S. Methodologically, attempts to fine tune the model are, according to Barnard et al. (1994), probably better preceded and accompanied by extensive data-analytic use of the model, to establish how well (or badly) it does, in fact, perform empirically.

As has been previously noted, much of the published literature claims that the application of the model has led to results that closely predict observed brand performance. This has only recently begun to be quantified systematically across many sets of data and reported (see for example, Ehrenberg 1999; Dekimpe et al. 1997; and Bhattacharaya 1997). Earlier work such as Uncles et al. (1994), report correlations for limited sets of data. This work has been previously summarised in § 2.5.1.

Ehrenberg (1988) summarises further necessary work as being:

(i) Prediction errors require investigation. Ehrenberg claims the errors are mostly small and may be of little marketing importance.

(ii) It would be helpful to develop closed-form approximations to the Dirichlet.

(iii) The lack of constancy of the parameter *S* should be investigated. As it stands, different values are obtained from different length time periods, and from different samples in the same market.

(iv) The issue of whether the two basic "diversity" parameters K and S are product class characterisics, and whether they are the same for different demographic sub-groups, different store-groups, different countries and so on, needs to be investigated.

(v) The meaning and interpretation of the *S* parameter should be clarified. Ehrenberg (1995b) goes on to ask "how far is the variation in S mere sampling? What other explanatory factors are involved? What effects does the variation in S have on the predictions of the model?", pg 130.

Ehrenberg and Uncles (1995) point out that while the Dirichlet model provides a good tool for handling stationary data, more work still needs to be done. Among the tasks they identify as important in the development of the Dirichlet is to investigate the variability of the parameter S. This task is taken up in this thesis.

CHAPTER FIVE

METHODOLOGY

5.1 Introduction

Previous chapters have identified a number of issues pertaining to the estimation of the Dirichlet model, and in particular, the estimation of the Dirichlet parameter *S*.

For example, Ehrenberg (1994) notes:

> that one needs to understand this variation (of S) and the meaning of S more generally. How far is the variation in S mere sampling? What other explanatory factors are involved? What effects does the variation in S estimates have on the predictions of the model?

Barnard et al. (1994) also note that an important outstanding issue is the variability and detailed meaning of the Dirichlet parameter S.

For the purpose of this thesis these issues are refined into the following objectives:

(1) To examine the sampling error of the Dirichlet *S* parameter.
(2) To examine the effect of sampling error on the Dirichlet predictions.

5.2 Method

The estimates of brand parameters used as inputs for the Dirichlet model are obtained from sample data. As such they are subject to sampling error. However, the procedure for estimating the Dirichlet predictions involves an iterative search procedure, and there is no analytical solution to the derivation of their standard errors. One way of obtaining the standard errors is simulate consumer purchasing records according to the distributions specified by the Dirichlet model, and compute the standard deviations of the estimates directly.

To obtain the standard error of the brand statistics of penetration and frequency simulated consumer panels were generated. In brief this involves generating a population of size n with Gamma distributed mean rates of buying u_n, and with purchase quantities Poisson distributed with mean u_n. From this, the brand statistics of penetration and frequency can be tabulated and when the procedure is repeated j times the empirical sampling distribution can be obtained.

The SAS system V6.12 was utilised to generate the Gamma and Poisson variates and in further detail the methodology for this section is provided below.

```
CODE                                                    COMMENT

%macro genpanel;
data test;
        do n = 1 to 10000;                              Generate n respondents
                 %do j = 1 %to 100;                     Generate j panels
                 /*Brim .010 2.1 */                     Brand parameters (Gamma
                                                        K, A)
                          rate = 2.1 * rangam(0,.010);  Generate respondent n's
                                                        rate of purchasing
                          purc&size = ranpoi(0,rate);   Generate respondent n's
                                                        purchases
                 %end;
                 output;
        end;
%mend looper;

%looper;                                                Invoke routine
run;
```

Once the simulated panels have been generated, the sampling distributions of the brand statistics of penetration and frequency are tabulated. This allows the computation of the Dirichlet parameter S for the upper and lower regions of the sampling distribution of the estimates of penetration and frequency for a given brand. Tabulations and summarising of the data was performed in a spreadsheet, and the fitting of the Dirichlet parameters, A, K, and S was performed using Dirichlet software; Kearns et al. (1998).

The original parameters used as inputs in the simulation study are detailed in Table 5.1. Simulation A contains a subset of three brands from the Ready-to-eat Cereal data, from Aske Research (1969). Simulation B is made up of eight hypothetical brands differing only in the observed average purchasing frequency, in this case ranging from 10 to 3. Simulation C uses data from Uncles and Ehrenberg (1990b).

Table 5.1
Simulation Parameters.

		Penetration	Frequency	*K*	A
Simulation A					
Category		.85	25.5		
Brand	C	.65	10.1	.353	18.6
	W	.43	8.5	.186	19.7
	R	.10	3.7	.049	7.5
Simulation B					
Category		.85	25.5		
Brand	A	.20	10	.064	31.09
	B	.20	9	.067	26.83
	C	.20	8	.070	22.70
	D	.20	7	.075	18.70
	E	.20	6	.081	14.86
	F	.20	5	.089	11.22
	G	.20	4	.103	7.80
	H	.20	3	.129	4.66
Simulation C: US Coffee Data					
Category		.31	5.0		
	Folgers	.11	3.2	.061	5.80
	Maxwell House	.10	3.3	.053	6.17
	Tasters Choice	.09	2.8	.055	4.61
	Other Brands	.08	3.0	.045	5.30
	Nescafe	.06	2.7	.037	4.43
	Sanka	.05	3.0	.03	5.50
	Maxim	.01	4.5	.010	4.50
	High Point	.01	2.6	.010	2.60
	Brim	.01	2.1	.010	2.10

CHAPTER SIX

RESULTS

6.1 Introduction

This chapter presents the results of the simulation experiments conducted to explore the effect of sampling error of the brand statistics on the estimation of the Dirichlet model, and primarily, the parameter *S*. The estimation of the Dirichlet S parameter is based on sample estimates of brand parameters, and these estimates are of course subject to variation due to sampling error. However, the relationship between the sampling distribution of these estimates and the sampling distribution of *S* is not available by any analytical method..

To obtain the sampling distribution of *S* and the predictions made by the Dirichlet model it is necessary to have the sampling distribution of the brand parameters of penetration and purchase frequency. To get the sampling distribution of brand penetration and purchase frequency, simulated panels were generated with the properties of the actual panel data.

The sampling distribution of the estimates of penetration and frequency is described in §6.2. The effect of sampling error of brand parameters on the Dirichlet parameter *S* is investigated in §6.3, and the effect on the Dirichlet model's predictions is examined in §6.4.

6.2 Parameters of the Sampling Distributions of the Simulated Panels

As an initial assessment of the validity of the simulation procedure the means of the distributions of the sampling distributions of the simulated panels are tabulated in Table 6.1.

The means of values of the original parameters for the simulated panels are close to the values obtained for the simulated samples. However the simulated samples are biased for purchase frequency for brands Maxim, High Point and Brim.

Table 6.1
Recovery of Simulation Parameters.

		Penetration				Frequency			
		Original	n=1000	n=5000	n=10000	Original	n=1000	n=5000	n=10000
Simulation A	C	.65	.653	.651	.653	10.1	10.1	10.0	10.01
Brands	W	.43	.433	.431	.433	8.5	8.5	8.5	8.44
	R	.10	.099	.100	.099	3.7	3.6	3.71	3.66
Simulation B	10	.20	.202	.199	.200	10	9.9	9.98	9.94
Brands	9	.20	.200	.199	.200	9	8.9	8.92	8.91
	8	.20	.199	.198	.199	8	7.9	7.96	7.91
	7	.20	.200	.201	.201	7	6.9	6.94	6.93
	6	.20	.202	.201	.201	6	5.9	5.95	5.95
	5	.20	.200	.199	.200	5	5.09	4.96	4.95
	4	.20	.200	.201	.202	4	3.92	3.98	3.97
	3	.20	.198	.201	.201	3	3.02	2.98	2.97
Simulation C:	Folgers	.11	.112	.110	.11	3.2	3.20	3.17	3.15
US Coffee	Maxwell House	.10	.096	.099	.10	3.3	3.23	3.26	3.25
Data	Tasters Choice	.09	.090	.090	.09	2.8	2.74	2.80	2.78
	Other Brands	.08	.080	.080	.08	3.0	2.92	2.98	2.96
	Nescafe	.06	.060	.060	.06	2.7	2.67	2.71	2.68
	Sanka	.05	.054	.055	.06	3.0	3.04	3.01	2.95
	Maxim	.01	.018	.018	.02	4.5	2.58	2.63	2.52
	High Point	.01	.014	.013	.01	2.6	1.99	2.01	1.89
	Brim	.01	.012	.012	.01	2.1	1.81	1.84	1.71

As a further check on the simulation process Table 6.2 compares the observed standard errors of *m*, the mean rate of buying with the theoretical value obtained from $\sqrt{m(1+m/k)/n}$, Ehrenberg (1988).

Table 6.2
Comparison of the standard error of *m* obtained from the simulated samples and the theoretical value $\sqrt{m(1+m/k)/n}$ [T=Theoretical, O=Observed].

	n=1000		n=5000		n=10000	
	T	O	T	O	T	O
Simulation A						
C	0.36	0.36	0.16	0.16	0.12	0.12
W	0.27	0.29	0.11	0.11	0.09	0.09
R	0.06	0.05	0.02	0.02	0.02	0.02
Simulation B						
10	0.24	0.24	0.11	0.11	0.08	0.08
9	0.23	0.23	0.1	0.1	0.07	0.07
8	0.21	0.21	0.09	0.09	0.06	0.06
7	0.17	0.17	0.07	0.07	0.05	0.05
6	0.14	0.14	0.06	0.06	0.04	0.04
5	0.1	0.1	0.05	0.05	0.03	0.03
4	0.09	0.09	0.04	0.04	0.03	0.03
3	0.06	0.06	0.03	0.03	0.02	0.02
Simulation C						
Folgers	0.04	0.04	0.02	0.02	0.02	0.02
Maxwell House	0.05	0.05	0.02	0.02	0.01	0.01
Tasters Choice	0.04	0.04	0.02	0.02	0.01	0.01
Other Brands	0.04	0.04	0.02	0.02	0.01	0.01
Nescafe	0.03	0.03	0.01	0.01	0.01	0.01
Sanka	0.03	0.03	0.01	0.01	0.01	0.01
Maxim	0.02	0.02	0.01	0.01	0.01	0.00
High Point	0.01	0.01	0.01	0.00	0.00	0.00
Brim	0.01	0.01	0.00	0.00	0.00	0.00
AVERAGE	0.11	0.11	0.05	0.05	0.03	0.03
$\sqrt{n} / \sqrt{1000}$			0.45		0.32	
$\frac{\text{Average s.e.}_n}{\text{Average s.e.}_{1000}}$			0.45		0.27	

The average standard error observed in the simulated samples and the theoretical value correspond, thus further analysis can proceed with confidence. One point to note is that the standard deviations of the simulated sampling distributions are a function of the sample size, but do not decline as a strict function of the square root of the sample size.

The Dirichlet model is conventionally estimated using estimates of brand penetration and frequency. In the case of market share estimates, brand penetration is estimated by:

$$\frac{\text{average no. of brands bought*proportion buying any brand}}{\text{no. of brands}} \qquad \text{(BUYER, 1990)}$$

This yields the same estimate of penetration for all brands. This means that for subsequent steps in the analysis the sampling distribution of brand penetration and purchase frequency is required rather than the mean rate of buying.

Table 6.3 sets out the computed standard deviations of the sampling distributions of both penetration and frequency for the simulated samples of n=1000, 5000, and 10,000 consumers.

Table 6.3
Standard Deviations of the Estimates of Penetration and Frequency from the Simulated Samples.

	Penetration			Frequency		
	n=1000	n=5000	n=10000	n=1000	n=5000	n=10000
Simulation A						
C	0.015	0.006	0.004	0.512	0.200	0.140
W	0.015	0.006	0.004	0.555	0.245	0.181
R	0.009	0.003	0.002	0.399	0.192	0.134
Simulation B						
10	0.012	0.006	0.004	1.081	0.505	0.394
9	0.014	0.005	0.004	0.957	0.414	0.270
8	0.013	0.006	0.004	0.894	0.408	0.280
7	0.012	0.006	0.004	0.702	0.303	0.219
6	0.011	0.006	0.004	0.566	0.281	0.200
5	0.012	0.006	0.004	0.448	0.205	0.163
4	0.011	0.006	0.004	0.353	0.149	0.010
3	0.013	0.005	0.004	0.212	0.097	0.072
Simulation C						
Folgers	0.010	0.005	0.004	0.331	0.162	0.109
Maxwell House	0.009	0.004	0.003	0.356	0.157	0.114
Tasters Choice	0.010	0.004	0.003	0.280	0.120	0.087
Other Brands	0.010	0.004	0.003	0.348	0.157	0.129
Nescafe	0.007	0.003	0.002	0.353	0.157	0.111
Sanka	0.007	0.003	0.002	0.456	0.211	0.141
Maxim	0.006	0.006	0.006	0.603	0.350	0.270
High Point	0.007	0.006	0.006	0.522	0.268	0.204
Brim	.007	0.006	0.006	0.459	0.232	0.198

As a further examination of these sampling distributions, Table 6.4 reports measures of Kurtosis and Skewness for the distribution as well as the One-Sample Kolmogorov-Smirnov statistics for testing whether the sampling distribution is normally distributed.

Only in the case of the brands Maxim, High Point and Brim are the sampling distributions of the parameters penetration and frequency different from a normal distribution with the observed mean and standard deviation.

Table 6.4
Kurtosis, Skewness and Kolmogorov-Smirnov Statistics from the Simulated Samples.

	Penetration				Frequency			
	Kurtosis	Skewness	Kolmogorov-Smirnov *Z*	2-tailed P	Kurtosis	Skewness	Kolmogorov-Smirnov *Z*	2-tailed P
Simulation A								
C	.33	-.06	.78	.57	-.31	.29	1.27	.08
W	-.31	-.07	.54	.93	.33	-.45	.64	.81
R	.18	-.58	.76	.61	.55	.31	.51	.98
Simulation B								
10	.98	-.25	.85	.47	-.35	.30	.45	.99
9	.28	-.35	.74	.64	-.26	.17	.86	.48
8	.07	.21	.56	.91	.39	-.07	.62	.83
7	.79	-.14	.60	.87	.32	-.26	.68	.71
6	-.12	-.01	.82	.51	.33	.12	.58	.89
5	-.33	-.01	.47	.98	.07	.01	.44	.99
4	-.19	-.35	.75	.63	-.23	.41	.75	.62
3	.04	-.16	.41	.99	.14	.08	.54	.94
Simulation C								
Folgers	.57	.34	.85	.47	-.66	-.24	.78	.58
Maxwell House	-.26	.35	.88	.42	-.29	-.13	.47	.98
Tasters Choice	1.25	.57	.61	.83	.85	.11	.71	.69
Other Brands	-.30	-.04	.75	.63	.02	-.25	.61	.85
Nescafe	.16	.23	1.06	.21	-.66	.14	.53	.94
Sanka	-.21	.24	.61	.85	-.24	.08	.77	.60
Maxim	21.6	-3.36	3.48	.00	-.29	-.13	1.13	.16
High Point	88.7	9.50	3.59	.00	22.66	-3.43	1.35	.05
Brim	87.5	9.40	3.64	.00	23.78	-3.48	1.30	.07

As illustrative examples the histograms and normal probability plots of the sampling distributions for brand C statistics are presented below in Figure 6.1.

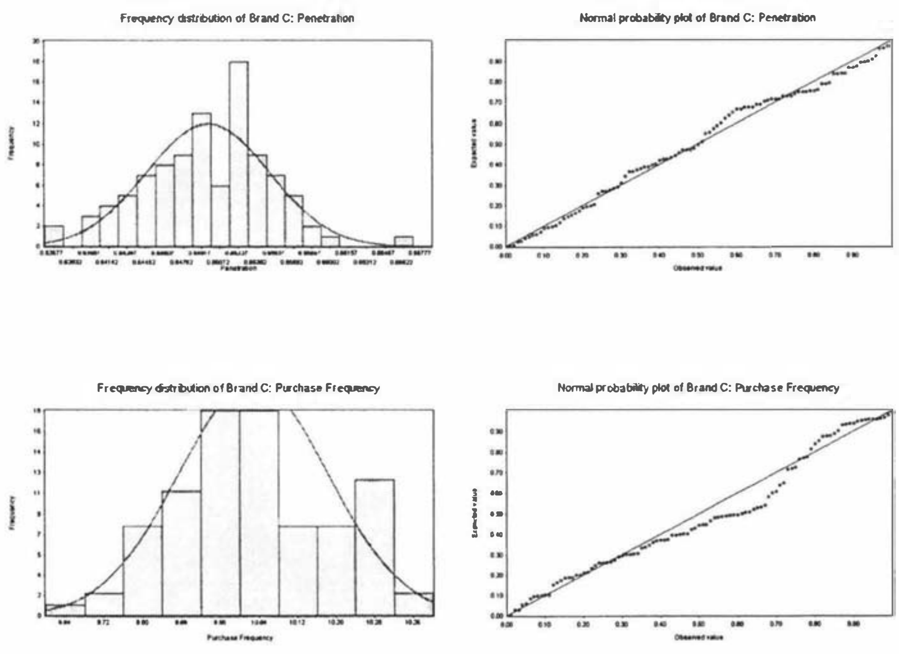

Figure 6.1
Histograms and Normal Probability Plots for Brand C Parameters

6.3 Estimates of the Dirichlet S Parameter

In this section the results from the previous section are used to perturb the original parameter estimates to investigate the effect of sampling error on estimation of the Dirichlet model. This procedure of perturbing the original parameters is necessary since no analytical solution is available relating the sampling error of the original parameters to the sampling error of the intermediate parameter *S*, nor the sampling error of the substantive outputs of the model such as the proportion buying once, twice and so on, or the proportion buying only that product.

To obtain a range of the input parameters of penetration and frequency for individual brands to represent sampling error, the standard deviations of the simulated sampling distributions as reported in Table 6.3 were used.

For example, from Tables 6.1 and 6.3, the parameters of Simulation A: Brand C: sample n=1000: were penetration .65 ($\sigma = .015$), and purchase frequency 10.1 ($\sigma = .512$). These results were used to perturb the original parameters to obtain two perturbed brands, representing the effect of sampling error. These perturbed brands were then used as inputs to estimate the Dirichlet parameter *S* for that set of parameters to give an indication of the variability of S due to sampling error.

Table 6.5 reports the range of the estimates of S obtained from the perturbed input parameters, as well as the mean value, which corresponds to the value of S obtained for the original parameter values.

Table 6.5
Mean and Range of Estimates of S

	S *[Original Parameters]*	Range of S from Perturbed Brand Parameters		
		n=1000	n=5000	n=10000
Simulation A				
C	3.02	0.51	0.19	0.14
W	1.90	0.82	0.37	0.28
R	5.76	5.81	2.56	1.76
Simulation B				
10	0.89	0.79	0.34	0.27
9	1.09	0.91	0.38	0.24
8	1.40	1.23	0.52	0.35
7	1.81	1.40	0.56	0.40
6	2.47	1.82	0.84	0.60
5	3.60	2.55	1.09	0.88
4	5.84	4.51	1.70	1.13
3	12.04	8.21	3.65	2.64
Simulation C				
Folgers	0.93	1.46	0.65	0.42
Maxwell House	0.78	1.36	0.51	0.37
Tasters Choice	1.36	1.88	0.72	0.52
Other Brands	1.02	1.72	0.68	0.57
Nescafe	1.31	2.77	0.97	0.66
Sanka	0.87	2.34	0.82	0.52
Maxim	0.11	0.46	0.25	0.20
High Point	1.26	5.31	1.54	1.07
Brim	2.74	24.50	3.60	2.82

These results for the simulated samples of consumer purchasing records show the variation in estimates of S from simulated Dirichlet samples. They show that the variation in S can be sometimes quite marked, and considerable variation in the estimates of S is observed for larger values of S, which are usually the result of smaller values of the input parameters.

To investigate the effect of sample size, Table 6.6 compares the ratio of the range in S for different sample sizes to that obtained for a sample of size n=1000.

Table 6.6
Range in estimates of S and Ratio of Range in S(n) to S (1000)

	Range in S for n=5000 / Range in S for n=1000	Range in S for n=10000 / Range in S for n=1000
Simulation A		
C	0.37	0.27
W	0.45	0.34
R	0.44	0.30
Simulation B		
10	0.43	0.34
9	0.42	0.26
8	0.42	0.28
7	0.40	0.29
6	0.46	0.33
5	0.43	0.35
4	0.38	0.25
3	0.44	0.32
Simulation C		
Folgers	0.45	0.29
Maxwell House	0.38	0.27
Tasters Choice	0.38	0.28
Other Brands	0.40	0.33
Nescafe	0.35	0.24
Sanka	0.35	0.22
Maxim	0.54	0.43
High Point	0.29	0.20
Brim	0.15	0.12
AVERAGE	0.40	0.29
$\sqrt{n} / \sqrt{1000}$	0.45	0.32

The range in S declines with increasing sample size. The comparison of the average ratio of the range of S for a given sample size to the range for sample size of n = 1000 shows that the range in estimates of S declines as a function of sample size, but somewhat faster than the $\sqrt{n}$.

The variability of *S* is further clarified in Figure 6.2 which depicts the relationship between different values of penetration and frequency and the computed values of *S* for a hypothetical market. This graph illustrates that there can be considerable variability in the estimates of S, illustrating the role of sampling error of the estimates of penetration and frequency on the resulting estimates of *S*.

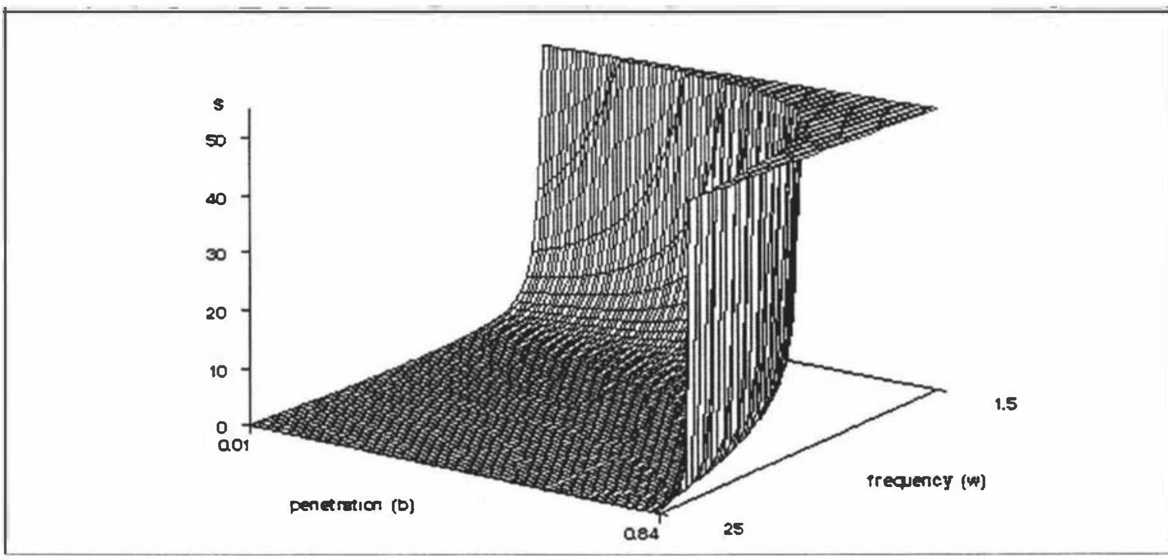

Figure 6.2
Simulated values of penetration and frequency and resulting estimate of *S*

The Relationship between the variability of *S* and the standard error of *m*, the mean rate of buying

The above analysis has demonstrated that there is considerable variation in the estimates of S, especially for small sample sizes. As an attempt to understand the relationship between the variation in S and underlying variation in the brand parameters from which *S* is estimated, Table 6.7 sets out the ratio of the range in *S* to the theoretical standard error of *m* the mean rate of buying.

Table 6.7
Ratio of range in estimates of *S* to the standard error of the mean rate of buying

	n=1000			n=5000			n=10000		
	s.e. of *m*[1]	Range in S	Range in S over s.e. of m	s.e. of *m*[1]	Range in S	Range in S over s.e. of m	s.e. of *m*[1]	Range in S	Range in S over s.e. of m
Simulation A									
C	0.36	0.51	1.42	0.16	0.19	1.19	0.11	0.14	1.24
W	0.28	0.82	2.98	0.12	0.37	3.01	0.09	0.28	3.26
R	0.15	5.81	37.97	0.07	2.56	37.65	0.05	1.76	36.67
Simulation B									
10	0.25	0.79	3.12	0.11	0.34	3.01	0.08	0.27	3.38
9	0.22	0.91	4.08	0.10	0.38	3.80	0.07	0.24	3.43
8	0.19	1.23	6.34	0.09	0.52	5.98	0.06	0.35	5.74
7	0.17	1.40	8.43	0.07	0.56	7.57	0.05	0.40	7.69
6	0.14	1.82	13.28	0.06	0.84	13.77	0.04	0.60	13.95
5	0.11	2.55	23.18	0.05	1.09	22.24	0.03	0.88	25.88
4	0.08	4.51	54.34	0.04	1.70	45.95	0.03	1.13	43.46
3	0.06	8.21	141.55	0.03	3.65	140.38	0.02	2.64	146.67
Simulation C									
Folger	0.04	1.46	36.50	0.02	0.65	30.95	0.02	0.42	28.00
Maxw	0.04	1.36	34.00	0.02	0.51	24.29	0.02	0.37	24.67
Tasters	0.03	1.88	62.67	0.02	0.72	45.00	0.01	0.52	47.27
Other	0.03	1.72	57.33	0.11	0.68	6.36	0.01	0.57	47.50
Nescaf	0.02	2.77	138.50	0.01	0.97	74.62	0.01	0.66	73.33
Sanka	0.03	2.34	78.00	0.01	0.82	63.08	0.01	0.52	57.78
Maxim	0.02	0.46	23.00	0.01	0.25	25.00	0.01	0.20	28.57
High	0.01	5.31	531.00	0.01	1.54	308.00	0.00	1.07	356.67
Brim	0.00	24.50	3066.00	0.00	3.60	1200.00	0.00	2.82	1410.00

[1] the standard error of m for a sample of size n being $\sqrt{m(1+m/k)/n}$, Ehrenberg (1988).

These results show that while the range of S to the theoretical standard error is approximately constant within a brand, the theoretical standard error of *m* does not account for the different ranges of *S*.

6.4 Effect of Sampling Error on Dirichlet Predictions

While the above analysis has demonstrated that there can be considerable variation in the estimate of S, the effects on the predictions of the model are at this stage unknown. This section explores the effect of the sampling error of the brand statistics on predictions made by the Dirichlet model.

Using the theoretical standard errors of m from Table 6.2, Tables 6.8, 6.9 and 6.10 give an indication of the effect of sampling error of the brand parameters on the Dirichlet predictions.

Table 6.8
Ratios of se *m* to Range of Dirichlet Predictions: Simulation A

					Proportion Buying		Purchase Frequency			100% Loyal		
	n	SE *m* [1]	Brand Share	Penetration	Once	Five +	of the brand	of the category	Share of Category	Proportion	Rate of Buying	Repeat Buying
C	1000	0.36	0.25	0.17	0.10	0.18	5.71	3.77	0.23	0.07	2.75	0.06
C	5000	0.16	0.22	0.15	0.09	0.16	4.99	3.32	0.20	0.06	2.40	0.05
C	10000	0.11	0.21	0.14	0.08	0.15	4.94	3.21	0.20	0.06	2.49	0.05
W	1000	0.27	0.25	0.22	0.16	0.26	8.08	5.89	0.29	0.11	8.43	0.09
W	5000	0.12	0.23	0.20	0.16	0.26	7.98	5.69	0.29	0.11	8.51	0.09
W	10000	0.09	0.24	0.18	0.16	0.26	8.33	5.85	0.30	0.11	9.05	0.09
R	1000	0.06	0.24	0.64	2.04	2.00	28.37	59.27	0.87	0.15	4.43	0.93
R	5000	0.03	0.22	0.48	2.09	2.13	30.52	61.70	0.93	0.16	4.68	0.97
R	10000	0.02	0.22	0.45	2.05	2.10	30.13	60.62	0.92	0.16	4.60	0.95
AVERAGE			0.23	0.29	0.77	0.83	14.34	23.26	0.47	0.11	5.26	0.36

Table 6.9
Ratios of se m to Range of Dirichlet Predictions: Simulation B

					Proportion Buying		Purchase Frequency			100% Loyal		
n		SE m [1]	Brand Share	Penetration	Once	Five +	of the brand	of the category	Share of Category	Proportion	Rate of Buying	Repeat Buying
10	1000	0.25	0.24	0.19	0.30	0.47	17.03	10.56	0.62	0.43	32.79	0.16
10	5000	0.11	0.26	0.21	0.31	0.48	17.79	10.89	0.65	0.44	35.75	0.16
10	10000	0.08	0.27	0.20	0.34	0.53	19.62	11.93	0.72	0.49	39.91	0.18
9	1000	0.22	0.26	0.25	0.34	0.51	17.09	11.78	0.61	0.34	30.11	0.18
9	5000	0.10	0.24	0.20	0.32	0.49	16.53	11.12	0.59	0.33	30.80	0.17
9	10000	0.07	0.23	0.23	0.29	0.45	15.25	10.33	0.55	0.30	28.30	0.16
8	1000	0.20	0.27	0.27	0.41	0.61	18.30	14.20	0.64	0.29	26.61	0.22
8	5000	0.09	0.27	0.27	0.41	0.61	18.68	14.23	0.65	0.29	28.04	0.22
8	10000	0.06	0.26	0.26	0.40	0.59	18.13	13.75	0.63	0.28	27.40	0.21
7	1000	0.17	0.25	0.29	0.45	0.64	16.92	15.21	0.58	0.20	17.38	0.23
7	5000	0.07	0.26	0.32	0.43	0.61	16.33	14.56	0.56	0.19	16.51	0.22
7	10000	0.05	0.25	0.30	0.43	0.62	16.69	14.78	0.57	0.19	16.96	0.23
6	1000	0.14	0.24	0.32	0.54	0.72	16.43	17.77	0.55	0.15	9.77	0.27
6	5000	0.06	0.28	0.39	0.59	0.79	18.24	19.54	0.61	0.16	10.36	0.30
6	10000	0.04	0.27	0.37	0.59	0.80	18.36	19.55	0.61	0.16	10.41	0.30
5	1000	0.11	0.25	0.43	0.70	0.87	16.20	22.41	0.53	0.11	4.75	0.35
5	5000	0.05	0.26	0.49	0.71	0.88	16.57	22.72	0.54	0.11	4.63	0.35
5	10000	0.04	0.28	0.46	0.78	0.98	18.64	25.23	0.60	0.13	5.26	0.39
4	1000	0.08	0.25	0.53	1.05	1.14	16.86	31.80	0.53	0.09	2.56	0.50
4	5000	0.04	0.27	0.64	0.99	1.08	15.91	30.03	0.50	0.08	2.31	0.47
4	10000	0.03	0.25	0.60	0.94	1.02	15.10	28.43	0.48	0.08	2.19	0.44
3	1000	0.06	0.26	0.89	1.57	1.30	14.56	43.09	0.44	0.06	1.20	0.68
3	5000	0.03	0.24	0.77	1.55	1.33	14.90	43.12	0.45	0.07	1.24	0.68
3	10000	0.02	0.26	0.87	1.64	1.40	15.64	45.43	0.47	0.07	1.29	0.72
Average			0.26	0.41	0.67	0.79	16.91	20.93	0.57	0.21	16.11	0.32

Table 6.10
Ratios of se M to Range of Dirichlet Predictions: Simulation B

	n	SE m [1]	Brand Share	Penetration	Proportion Buying: Once	Proportion Buying: Five +	Purchase Frequency: of the brand	Purchase Frequency: of the category	Share of Category	100% Loyal: Proportion	100% Loyal: Rate of Buying	Repeat Buying
Folgers	1000	0.05	3.69	0.86	2.59	2.37	27.10	28.56	6.20	6.98	40.10	2.28
Folgers	5000	0.02	3.94	0.89	2.77	2.57	29.59	30.47	6.74	7.66	46.09	2.42
Folgers	10000	0.02	3.92	0.93	2.65	2.46	28.34	29.20	6.46	7.32	44.27	2.32
Maxwell	1000	0.05	3.37	0.70	2.73	2.52	29.12	30.01	6.81	7.89	43.49	2.39
Maxwell	5000	0.02	3.52	0.78	2.64	2.47	28.71	29.04	6.68	7.81	45.00	2.31
Maxwell	10000	0.02	3.58	0.79	2.69	2.53	29.43	29.67	6.84	8.02	46.45	2.36
Tasters	1000	0.04	3.63	1.05	3.38	2.77	29.90	37.18	6.31	6.19	38.65	3.06
Tasters	5000	0.02	3.38	0.95	3.15	2.64	28.61	34.64	6.02	5.90	38.51	2.83
Tasters	10000	0.01	3.50	0.99	3.23	2.71	29.41	35.54	6.18	6.06	39.71	2.90
OtherB	1000	0.04	3.83	1.03	3.78	3.21	35.72	41.66	7.85	8.36	49.79	3.40
OtherB	5000	0.02	3.68	0.94	3.70	3.22	36.09	40.74	7.89	8.47	53.57	3.30
OtherB	10000	0.01	3.75	0.83	4.23	3.72	41.75	46.59	9.09	9.85	63.04	3.77
Nescafe	1000	0.03	3.44	0.91	5.81	4.47	47.82	63.89	9.83	9.59	59.65	5.36
Nescafe	5000	0.01	3.57	0.99	5.55	4.42	47.51	61.06	9.73	9.46	62.78	5.07
Nescafe	10000	0.01	3.70	1.07	5.55	4.43	47.63	61.05	9.76	9.46	63.27	5.07
Sanka	1000	0.03	3.65	0.87	6.71	5.47	60.85	73.81	13.39	14.37	81.78	6.11
Sanka	5000	0.01	3.97	1.01	6.61	5.60	62.82	72.69	13.72	14.91	92.63	5.96
Sanka	10000	0.01	3.89	1.02	6.24	5.31	59.65	68.64	13.01	14.16	89.15	5.62
Maxim	1000	0.02	5.69	1.85	8.96	8.52	104.50	98.60	30.41	35.72	123.75	7.86
Maxim	5000	0.01	10.80	3.44	13.28	12.96	161.39	146.23	50.01	58.40	175.23	11.57
Maxim	10000	0.00	14.42	4.58	15.98	15.69	196.15	175.88	62.05	72.10	206.68	13.89
High Point	1000	0.01	6.18	2.85	30.95	20.15	215.70	339.85	43.57	41.89	244.36	30.04
High Point	5000	0.00	10.61	5.38	32.37	23.20	247.74	355.77	49.76	47.41	308.22	30.59
High Point	10000	0.00	14.48	7.60	34.71	25.03	267.33	381.50	53.75	50.88	335.54	32.75
Average			5.34	1.76	8.76	7.02	78.87	96.34	18.42	19.54	99.65	8.05

These results, however, do not yield tractable results. There are wide variations in the relationship in the range of predictions to the standard error of the input parameters, rather than being proportionate to the standard error of *m*, the mean rate of buying.

CHAPTER SEVEN

DISCUSSION

7.1 Introduction

The Dirichlet model is a stochastic approach to the modelling of consumer behaviour. By making probabilistic assumptions about purchase incidence and brand choice the model specifies the probability of a consumer making n purchases in a time period and which brand is bought for each of the n purchases. By aggregating these probabilities over consumers the model is capable of describing and predicting a number of aspects of aggregate consumer behaviour of interest to marketers. Within the Dirichlet model, a Poisson process describes purchase incidence for an individual consumer, and the long run rates of purchasing incidence is distributed gamma across consumers. Brand choice is represented by multinomial probabilities for different brands in the market, and consumers differ from each other in the size and composition of their brand repertoires and in their brand choice probabilities within these repertoires according to a Dirichlet distribution.

The Dirichlet model integrates the reported regularities, and predicts many aggregate brand performance measures. These measures are the distribution of purchases for a brand, the proportion of a brand's buyers buying that brand only, and the proportion of people purchasing a brand, given that they have previously purchased that brand. When these predictions are compared with observed figures, Ehrenberg claims that it is not unreasonable to expect to obtain correlations in the order of 0.9 and sometimes much higher, (Ehrenberg 1975, Ehrenberg and Bound 1993).

The Dirichlet predictions are obtained from an iterative minimisation algorithm and their sampling properties are unknown.

7.2 Background

This thesis has investigated sampling issues in the estimation of the Dirichlet model of consumer behaviour. The motivation for this thesis was provided by Ehrenberg (1995b) who identified the need to clarify the meaning and interpretation of the S parameter. Ehrenberg (1995b) goes on to ask "how far is the variation in S mere sampling? What other explanatory factors are involved? What effects does the variation in S have on the predictions of the model?", pg 130.

As a development of Ehrenberg's suggestion this thesis has sought in particular to study the effect of sampling error on the estimation of the Dirichlet parameter *S* and predictions from the model. This has been achieved by simulating consumer panels of brand purchases with known parameters distributed according to the assumptions of the Dirichlet model to obtain information about the sampling distribution of the original parameters, then using this information to perturb the original parameters and then estimating the parameter S, thus obtaining values of *S* for the sampling distribution of the original brand parameters.

If it were to turn out that the sampling errors were a readily identifiable function of the known sampling error of the brand parameter of mean rate of buying then this would provide the analyst with useful information in the application of the model. Alternatively modifications to existing software solutions could adopt the a simulation procedure to directly estimate the effect of sampling error.

7.3 Results

Goodhardt et al. (1984) identify that one particular aspect of the estimation of *S* that needs further study is the constancy of *S* over different time periods. They note that the parameter *S* in the model should be constant irrespective of the length of time period considered, but it varied from 1.0 when estimated for 4 week data to 2.2 for the 12 week data, to 1.8 in the year, in a study of the buying of toothpaste brands.

Within this thesis this theme is taken up by investigating the effect of sampling error of the input parameters on the esimtates of the Dirichlet model.

The variation of *S* for the sampling distribution of the original parameters shows at times considerable variation in the values obtained for *S*. In principle, *S* varies to a larger degree for larger values of *S*, but this is not a tractable solution since one might require knowledge independent of the computed value of *S*

The relationship with thestandard error of m, the mean rate of buying, $\sqrt{m(1+m/k)/n}$, Ehrenberg (1988), was then investigated. This however does not yield a consistent factor. For example in Simulation A, the factor ranges from 1.24 for Brand C to 36.67 for Brand R.

Further clarification of the role of sampling error of the brand parameters of penetration and frequency is provided by Figure 6.1. This illustrates the effect of different values of penetration and frequency on the estimated value of S. Since the response surface is not linear, the magnitude of the effect of sampling error depends on the slope of the response surface around the original parameters of penetration and frequency.

A separate issue is the effect of sampling error on predictions obtained from Dirichlet predictions. The perturbed brand parameters were used to obtain estimates of the range of predictions for a number of characteristics predicted by the Dirichlet model, such as the proportions buying times, Share of category requirements, and the proportion who only buy that brand. Again no consistent relationship was found with the standard error of m, the mean rate of buying.

Since a function relating the sampling distribution of a brands' input parameters to the sampling distribution of the Dirichlet parameter S, or the predictions of aspects of brand performance was not identified, in practice a simulation solution will be required. One suggested approach to this would be to use the analytical solution to the sampling error of m, the mean rate of buying. This result could then be used to generate the sampling distribution of m for the brand, and separate Dirichlet models fitted to the perturbed brand parameters to investigate the sampling distribution of the parameters. This would then allow an analyst to investigate whether differences in estimates of S were attributable to sampling error or perhaps evidence of violations of the assumptions of the model.

References

Aaker, D A (1991) *Managing brand equity*, New York: The Free Press.

Ambler, T and C Styles (1995) Brand equity: Towards measures that matter, *London Business School Working Paper Series: 95-902*

Anon, (1993) What's in store for partnerships, *Discount Merchandiser*, 10, 10-13.

Anon, (1993) Price cuts, coupons rule butter, margarine sales, *Brandweek*, 34, 38, 54-56.

Anon, (1993) Research: survey reveals fall in brand loyalty, *Marketing*, July 8, 7.

Anon, (1993) Brand loyalty: Can it be revived? *Industry Week*, 242, 7, 11-14.

Anscombe, F J (1950) Sampling theory of negative binomial and logarithmic series distributions, *Biometrika*, 37, 358-382.

Aske Research (1969) *Loyalty reports: Ready to Eat Cereals:* Aske Research: London

Baldinger, A L and J Rubinson (1996) Brand loyalty: The link between attitude and behaviour, *Journal of Advertising Research,* 36, 6, 22-34.

Barnard, N R and A S C Ehrenberg (1990) Robust measures of consumer brand beliefs, *Journal of Marketing Research,* 27, November, 477-84.

Barnard, N R and A S C Ehrenberg (1997) Advertising: Strongly persuasive or nudging *Journal of Advertising Research,* 37, 1-31.

Barnard, N R, A S C Ehrenberg, K Hammond and M D Uncles (1994) Some debatable Dirichlet discrepancies, *London Business School Working Paper Series.*

Barnard, N R, A S C Ehrenberg, and J Scriven (1998) Branding and values, *Admap,* June, 47-49.

Bass, F M (1974) The theory of stochastic preference and brand switching, *Journal of Marketing Research,* 11, February, 1-20.

Bass, F M, A Jeuland and G P Wright (1976) Equilibrium stochastic choice and market penetration theories: Derivations and comparisons, *Management Science,* 1051-1065.

Bass, F M, M Givon, M U Kalwani, D Reibstein and G P Wright (1984) An investigation into the order of the brand choice process, *Marketing* Science, 267-287.

Bass, F M; E A Pessemier and D R Lehmann (1972) An experimental study of the relationships between attitudes, brand preference and choice, *Behavioural Science*, 17, November, 532-541.

Bass, F M and I Pilon (1980) A stochastic brand choice framework for econometric modelling of time series market share behaviour, *Journal of Marketing Research,* 17, November, 486-497.

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

Bemmaor, A (1995) Commentary, *in Research Traditions in Marketing,* G Laurent, G Lilien and B Pras ed., 109-115.

Bhattacharya, C B (1997) Is your brand's loyalty too much, too little, or just right?: Explaining deviations in loyalty from the Dirichlet norm, *International Journal of Research in Marketing,* 14, 421-435.

Bhattacharya, C B, P S Fader, L M Lodish and W DeSarbo (1996) The relationship between the marketing mix and share of category requirements, *Marketing Letters,* 7:1, 5-18.

Blattberg, R C (1981) Evaluation of stochastic brand choice models, in *Marketing Decision models,* eds. R L Schultz and A A Zoltners, New York: North Holland, 183-206.

Blattberg, R C and S K Sen (1976) Market segments and stochastic brand choice models, *Journal of Marketing Research,* 13, February, 34-45.

Castleberry, S, N R Barnard, T P Barwise, A S C Ehrenberg, and F Dall-Olmo Riley (1993) Individual attitude variation over time, *Journal of Marketing Management*, 9, 314-321.

Chatfield, C, A S C Ehrenberg and G J Goodhardt (1966) Progress on a simplified model of stationary purchasing behaviour, *Journal of the Royal Statistical Society, Series A,* 129, 317-367.

Chatfield, C and G J Goodhardt (1973) A consumer purchasing model with Erlang interpurchase times, *Journal of the American Statistical Association,* 68, 344, 828-835.

Chatfield, C and G J Goodhardt (1975) Results concerning brand choice, *Journal of Marketing Research,* 12, 110-113.

Colombo, R A and D G Morrison (1989) A brand switching model with implications for marketing strategies, *Marketing Science*, 8, 89-99.

Cooper, L G and M Nakanishi (1988) *Market Share Analysis,* Boston: Kluwer.

Corlett, W J (1966) A comment on progress on a simplified model of stationary purchasing behaviour, in Chatfield, C, A S C Ehrenberg and G J Goodhardt (1966) Progress on a simplified model of stationary purchasing behaviour, *Journal of the Royal Statistical: Series A,* 129, 317-367.

Dalal, S R, J C Lee and D J Sabavala (1984) Prediction of individual buying behaviour: A Poisson-Bernoulli model with arbitrary heterogeneity, *Marketing Science,* 3, 4 (Fall) 352-366.

Dall'Olmo Riley, F, A S C Ehrenberg, S B Castleberry, T P Barwise, N R Barnard, (1997)The variability of attitudinal repeat rates, *International Journal of Research in Marketing,* 14, 437-450.

Dekimpe, M G and D M Hanssens (1995) Empirical generalisations about market evolution and stationarity, *Marketing Science,* 14, 3, 2, G109-G121.

Dekimpe, M G, J E M Steenkamp, M Mellens and P Vanden Abeele (1997) Decline and variability in brand loyalty, *International Journal of Research in Marketing,* 14, 405-420.

Dodson, J A, A M Tybout and B Sternthal (1978) Impact of deals and deal retraction on brand switching, *Journal of Marketing Research,* 15(1) 72-81.

Dunn, R, S Reader and N Wrigley (1983) An investigation of the assumptions of the NBD model as applied to purchasing at individual stores, *Applied Statistics*, 32, 3, 249-259.

East, R and K Hammond (1996) The erosion of repeat purchase loyalty, *Marketing Letters,* 7, 2, 163-171.

Ehrenberg, A S C (1959) The pattern of consumer purchases, *Applied Statistics,* 8, 26-41.

Ehrenberg, A S C (1965) An appraisal of Markov brand switching models, *Journal of Marketing Research,* II, November, 347-62.

Ehrenberg, A S C (1975) *Data Reduction,* John Wiley and Sons, New York.

Ehrenberg, A S C (1988) *Repeat Buying: Theory and Applications,* London: Charles-Griffin.

Ehrenberg, A S C (1995a) Theory or well based results: which comes first?, in *Research Traditions in Marketing,* G Laurent, G Lilien and B. Pras ed., 79-105.

Ehrenberg, A S C (1995b) Author's reply by A S C Ehrenberg, in *Research Traditions in Marketing,* G Laurent, G Lilien and B. Pras ed., 128-131.

Ehrenberg, A S C (1996) Towards an integrated theory of consumer behaviour, *Journal of the Market Research Society,* 38, 4, October, 395-427.

Ehrenberg, A S C, K Hammond and G J Goodhardt (1994) The after-effects of price-related consumer promotions, *Journal of Advertising Research,* 34, July/August, 1-10.

Ehrenberg, A S C, G J Goodhardt and T P Barwise (1990) Double jeopardy revisited, *Journal of Marketing,* Vol 54, July, 82-91.

Ehrenberg, A S C (1993) If you're so strong why aren't you bigger?, *Admap,* October, 13-14.

Ehrenberg, A S C (1997) Description and prescription, *Journal of Advertising Research,* 37, November/December, pg 17-22.

Ehrenberg, A S C and J A Bound (1993) Prediction and predictability, *Journal of the Royal Statistical Society: Series A.,* 156, part 2, 167-206.

Ehrenberg, A S C and M D Uncles (1995) Implications of Dirichlet-type markets, *Working Paper*, London Business School.

Ehrenberg, A S C and M D Uncles (1997) Dirichlet-type markets: A position paper, *South Bank University and University of New South Wales Working Paper.*

Engel, J F, R D Blackwell and P W Miniard (1995) *Consumer Behaviour*, 8th ed., Dryden Press, Orlando, Florida.

Fader, P S and James M Lattin (1993) Accounting for heterogeneity and non-stationarity in a cross sectional model of consumer purchase behaviour, *Marketing Science,* 12, 3, 304-317.

Fader, P S and D C Schmittlein (1993) Excess behavioural loyalty for high share brands: deviations from the Dirichlet model for Repeat Purchasing, *Journal of Marketing Research,* 30 , November, 478-493.

Fournier, S and J L Yao (1997) Reviving brand loyalty: A reconceptulisation within the framework of consumer-brand relationships, *International Journal of Research in Marketing,* 14, 451-472.

Frank, R E (1962) Brand choice as a probability process, *Journal of Business,* 35, January, 43-56.

Givon, M and D Horsky (1978) Market share models as approximators of aggregated heterogeneous brand choice behaviour, *Management Science,* 24, 13, 1404-1416.

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

Grahn, G L (1969) NBD model of repeat purchase loyalty: An empirical investigation, *Journal of Marketing Research,* VI, February, 72-78.

Greene, J D (1989) Commentary on a brand switching model with implications for marketing strategies, *Marketing Science,* 8, 1, Winter, 104-105.

Greenwood, M and G U Yule (1920) An enquiry into the nature of the frequency distributions representative of multiple happenings, with particular reference to the occurrence of multiple attacks of disease or repeated accidents, *Journal of the Royal Statistical Society, Series A,* 83, 255-279.

Hariharan, H S and R P Velu (1993) On estimating Dirichlet parameters - a comparison of initial values, *Journal of Statistical Computational Simulation,* 48, 47-58.

Hauser, J R and B Wernerfelt (1990) An evaluation cost model of consideration sets, *Journal of Consumer Research,* 16, 393-408.

Herniter, J (1971) A probabilistic market model of purchase timing and brand selection, *Management Science*, 18, 102-113.

Howard, J A and J Sheth (1969) *The theory of buyer behaviour*, New York: John Wiley and Sons.

Hutchinson, P and L Marchant (1998) Added vlaies or propensities to buy?: Beyond the Dirichlet, *Admap,* May, 31-33.

Jacoby, J and R W Chestnut (1978) *Brand Loyalty: Measurement and Management*, Wiley, New York.

Johnson, N L and S Kotz (1969) *Discrete Distributions*, John Wiley and Sons, New York.

Jones (1973) A composite heterogenous model for brand choice behaviour, *Management Science,* 19, January, 499-509.

Jones, J M (1979) Whither stochastic choice models, *Working paper*, Graduate School of Management, University of California, Los Angeles.

Kahn, B E; M U Kalwani and D G Morrison (1988) Niching versus change-of-pace brands: Using purchase frequencies and penetration rates to infer brand positionings, *Journal of Marketing Research*, 25, 384-390.

Kamkura, W A and G J Russell (1989) A probabilistic choice model for market segmentation and elasticity structure, *Journal of Marketing* Research, 26, November, 379-90.

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

Khan, B and D G Morrison (1989) A note on 'random purchasing': Additional insights from Dunn, Reader and Wrigley, *Applied Statistics,* 38, 1, 111-114.

Krishnamurthi, L and S P Raj (1991) An empirical analysis of the relationship between brand loyalty and consumer price elasticity, *Marketing Science*, 10, Spring, 172-183.

Lal, Rajiv. Padmanabhan, V (1995) Competitive response and equilibria. *Marketing Science* 14(3(Part 2)): G101-G108.

Lattin, J M (1990) Measuring preference from scanner panel data: Filtering out the effects of price and *promotion, Unpublished working paper,* Graduate School of Business, Stanford University.

Lenk, P J, A R Rao and V Tibrewala (1993) Non stationary conditional trend analysis: An application to scanner panel data, *Journal of Marketing Research*, 30, 3, August, 288-304.

Light, L (1998) Brand loyalty management - the new marketing basic, *Admap,* May, 28-30.

Lilien, G; P Kotler and S K Moorthy (1992) *Marketing Models*, Prentice Hall, Englewood Cliffs, New Jersey

Little, J D C and E Anderson (1994) A product choice model with marketing, filtering and purchase feedback, *MIT Sloan School of Management Working Paper 3670-94.*

Massy, W F (1966) Order and homogeneity of family specific brand switching processes, *Journal of Marketing Research,* 3, February, 48-54.

Massy, W F, D B Montgomery and D G Morrison (1970) *Stochastic models of buyer behaviour,* Cambridge, MA:MIT Press.

McPhee, M (1963) *Formal Theories of Mass Behaviour,* Free Press, Glencoe, New York.

Montgomery, D B and A B Ryan (1974) Stochastic models of consumer choice behaviour. In *Consumer Behaviour: Theoretical Foundations,* eds. T S Robertson, and J S Ward.

Morrison, D G (1969) Conditional Trend Analysis: A model that allows for nonusers, *Journal of Marketing* Research, VI, August, 342-6.

Morrison, D G (1969) A series approximation for Negative Binomial Parameter estimation, *Journal of Marketing* Research, VI, August, 355-6.

Morrison, D G and D C Schmittlein (1988) Generalizing the NBD model for customer purchases: What are the implications and is it worth the effort, *Journal of Business and Economic* Statistics, April, 6, 2, 145-159.

Morrison, D G and D C Schmittlein (1981) Predicting future random events based on past performance, *Management* Science, 9 (September) 1006-1023.

Morrison, Donald G. Silva-Risso, Jorge (1995) A latent look at empirical generalizations, *Marketing Science*, 14(3(Part2)): G61-G70.

Mosimann, J E (1962) On the compound multinomial distirbution, the multivariate B-distribution, and correlations between proportions, *Biometrika,* 49, 65-82.

Mosimann, J E (1984) Size and shape analysis. In *Encyclopaedia of the Statistical Sciences,* eds. S Kotz and N L Johnson, New York: Wiley.

Murthi, B P S, K Srinivasan, P R Tadikamalla (1993) Robustness of the negative binomial distribution to model purchase incidence: A simulation study, *Communications in Statistics: Simulation*, 22 (2) 399-420.

Narayanan, A (1991) Small sample properties of parameter estimates in the Dirichlet Distribution, *Communications in Statistics-Simulation,* Vol 20, 647-666.

Ortmeyer, G; J M Lattin and D B Montgomery (1991) Individual differences in response to consumer promotions, *International Journal of Research in Marketing,* 8, September, 169-186.

Peter, J P and J C Olson (1995) *Understanding Consumer Behaviour*, Irwin: Illinois

Pfouts, R W (1994) Five recent econometric revolutions, *Atlantic Economic Journal,* 22(1) 22-25.

Rosenberg, I J and J A Czepiel (1983) A marketing approach to customer retention, *Journal of Consumer Marketing,* 2, 45-51.

Russell, G J (1994) Averaging out the past: When simple loyalty measures work, *Marketing Science Conference Proceedings*, University of Arizona.

Schmittlein, D C, A C Bemmaor and DG Morrison (1985) Why does the NBD model work? Robustness in representing product purchases, brand purchases and imperfectly recorded purchases, *Marketing Science,* 4, 3 (Summer), 255-266.

Schmittlein, D C, and DG Morrison (1983) Prediction of future random events with the condensed Negative Binomial Distribution. *Journal of the American Statistical Association,* 78, 382, 449-457.

Schmittlein, D C, L G Cooper and D G Morrison (1993) Truth in concentration in the land of (80/20) Laws, *Marketing Science*, 12(2), 167-193.

Sharp, B and A Sharp (1997) , Loyalty programs and their impact on repeat-purchase loyalty patterns, *International Journal of Research in Marketing*, 14, 473-486.

Shoemaker, R W; R Staelin, J B Kadane and F R Shoaf (1977) Relation of brand choice to purchase frequency, *Journal of Marketing Research,* 14, 458-468.

Srinivasan, V and T Kibarian (1989) Purchase event feedback: Fact or fiction? *Marketing Science Conference Proceedings*, Fuqua School of Business, Duke University.

Stern, P. (1994) Patterns of store pharmaceutical prescribing, *PhD Thesis,* London Business School.

Stern, P (1997) Long run loyalty: An exploratory study, *Marketing and Research Today,* May, 70-78.

Strang, R A (1975) The relationship between advertising and promotion in brand strategy, *Marketing Science Institute,* Cambridge.

Uncles, M D (1989) *BUYER*: Buyer behaviour software, CMAC Report, London Business School.

Uncles, M D and A S C Ehrenberg (1990a) Industrial buying behaviour: Aviation fuel contracts, *International Journal of Research in Marketing*, (7), 57-68.

Uncles, M D and A S C Ehrenberg (1990b) The buying of packaged goods at US retail chains, *Journal of Retailing*, 66(3), 278-296.

Uncles, M D; K A Hammond, A S C Ehrenberg and R E Davis (1994) A replication study of two brand-loyalty measures, *European Journal of Operational Research,* 76, 375-384.

Uncles, M D and A S C Ehrenberg and K Hammond (1995) Patterns of Buyer Behaviour: regularities, models and extensions, *Marketing Science*, 14, 3, Part 2, G71-G78.

Uncles, M D and G Laurent (1997) Editorial, *International Journal of Research in Marketing,* 14, 399-404.

Wheat, R D and D G Morrison (1990) Estimating purchase regularity with two interpurchase times, *Journal of Marketing Research,* 27, February, 87-93.

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