PART VI

R & D

CHAPTER 11

THE GROWTH OF THEORY

11.1. Research into Buyer Behaviour

This chapter is concerned with more general questions of theory, and Chapter 12 with more general problems of practical applications. The new Chapter 13 on the Dirichlet integrates the findings more fully.

The results concerning repeat-buying and multi-brand buying in the earlier parts have been described with little reference to *other* research into buyer behaviour. General methodological problems of model-building are discussed in the later sections of the chapter – for example, how the various findings in the earlier parts of the book arose, how they inter-relate, and how they help to explain each other. But in §§11.2 and 11.3 we first comment briefly on the general background of other theoretical and empirical work in this area.

Traditionally, the task of understanding consumer behaviour seems to have been approached along two parallel lines, the "practical" one of collecting facts, and the "academic" one putting forward theories. These two activities have seldom met. Thus various theories of consumer behaviour have been developed with little recourse to organised facts (as is summarised in § 11.2), and a great variety of facts have been collected with little if any conceptualisation (as discussed in § 11.3).

The great need therefore arose to let theory meet with fact, and vice versa. The kind of work described in this book has in fact centred on first of all uncovering regular patterns in empirical data, and then synthesising "theory", by first modelling, and then inter-relating, these regularities. Such work typically starts by examining some factual data. The data must be judged to be "relevant" – here records of *buying* behaviour – but apart from this there need be no very explicit theoretical concepts or any *a priori* framework of hypotheses.

The immediate results of such an examination of data are then a variety of separate relationships, each describing some more or less isolated aspect of the system. This is the lowest level of model-building, and there are two main criteria for assessing success. One is the degree of generalisation achieved: has the empirical regularity been shown to hold under a wide enough range of conditions to make it worth exploring further? The second criterion concerns the simplicity of each result: is it simple enough to facilitate its further use?

As regards the findings in this book, their *generality* (and its limitations) has already been extensively stressed in the earlier parts and need not be enlarged on further. The *simplicity* of the findings will also be fairly self-evident. It derives partly from the fact that of the vast number of variables which might have mattered, only two did (the penetration b and the average purchase frequency w), and partly from the virtual absence of numerical coefficients in the models (none in the repeat-buying theory, and only two or three in the as yet less developed multi-brand buying area of Part V).

Once a number of such simple and lawlike regularities have been uncovered, they help to suggest hypotheses for further work. Some of these hypotheses may be directed simply at the search for further low-level empirical regularities, and the interplay between theory and fact at this level is discussed in § 11.4. Other hypotheses arising from previous results may be more ambitious, and are illustrated in § 11.5. Thus whilst the repeat-buying results described in Parts I to IV and the multi-brand buying results in Part V are in essence *descriptive* at a fairly straightforward level (they tend to summarise how one aspect of buyer behaviour is related to another), in § 11.5 we discuss how such findings have started to inter-connect more widely, and how this is beginning to produce some of the "reasons why" for the basic empirical findings themselves.

The theoretical structure developed so far tends to work well, but there are also discrepancies and gaps which still need to be pursued. These are discussed further in § § 11.6 and 11.7.

11.2. Models Without Facts

The academic tendency in consumer research has been to try to explain why consumers behave as they do, without knowing or understanding in any quantitative detail just how they *do* in fact behave.

One of the most developed examples of this "explanatory" approach is probably the Howard-Sheth theory of buyer behaviour. This links up in verbal terms constructs such as "Attention", "Attitude", "Motive", "Comprehension", "Intention", marketing inputs such as "Price", "Quality", "Availability", social variables such as "Family", and "Reference Groups", and aspects of actual behaviour such as "Overt Search" and "Purchase", all this leading to "Satisfaction" and other forms of feedback. The model is illustrated in Fig. 11.1 overleaf which is taken from the authors' text [Howard and Sheth 1969].

The main problem here seems to be that there is no detailed evidence of just how any of these ideas relate to any specific facts of consumer behaviour. Some of these facts are measured in commercial market research but little quantitative use of them appears to be made. Instead, it seems to be implied that *new* kinds of data are needed in order to put some flesh on the theoretical structure. Indeed, even when the Howard-Sheth model was exposed to a factual test with specially collected data [Farley and Ring 1970], the conclusions were that "the test put extreme pressure on the data" (not on the *model*!), and that "considerably improved data-collection techniques and procedures will be needed before the full empirical potential of such models will be realised". It remains unclear whether or not the Howard-Sheth type of model is consistent with any of the empirical regularities of buyer behaviour which have been described in this book. It is not even very clear how to set about trying to establish this. One difficulty is the inherent lack of operational definitions of the concepts used.

The Howard-Sheth theory is only one of quite a number of more or less intuitive theories of consumer behaviour, some verbal and some more mathematical. Largely verbal theories of buyer behaviour developed from the marketing or behavioural point of view are covered in texts and readings on marketing and consumer behaviour (see for instance Engel et al. 1968, Ehrenberg and Pyatt 1970, Kollat et al. 1970, Robertson 1970, Walters and Paul 1970, Holloway et al. 1971), a succinct account of several being for example given in Kotler's well-known *Marketing Management* [1967, Chapter 4]. In addition, there are of course the various more or less classical theories of the consumer in economics, as reviewed for instance by Ferber [1962]. Examples of more mathematical formulations of consumer decision-processes are those of Nicosia [1966] using explicit formulae, and the computer-based simulation procedures best known about in the marketing area probably in the form of the work of Amstutz [1967]. More generally, the quantitative "management science" approach to marketing problems is covered in texts such as King [1967], Montgomery and Urban [1969, 1970], and Simon and Freimer [1970]. Stripped of the mathematics and the assumptions however, it does not appear that such theories contain any generalisable knowledge of consumer behaviour, or that the analytic techniques used have yet proved themselves in routine use.

From the narrow point of view of empirical buyer behaviour in the sense of purchase frequency and brand-choice as examined in this book, only one aspect of these various theories seems relevant at any

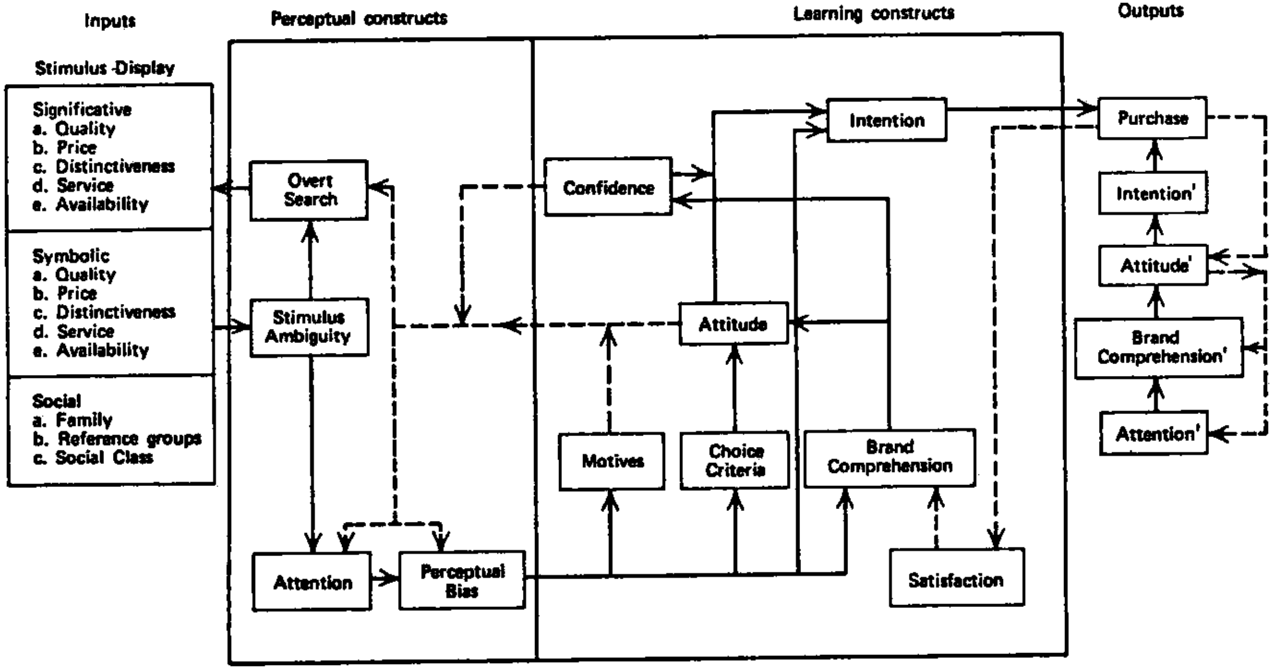

Fig. 11.1 A simplified description of the Howard-Sheth theory of buyer behaviour. Solid lines indicate flow of information; dashed lines, feedback effects. [From Howard and Sheath 1969.]

immediately actionable level. It is that none of the theories appear either to lead towards, or to provide a test for, such systematic findings as have been reported here. For example, none of the theories seem to require that the structure of repeat-buying and brand-choice must vary from brand to brand or from product to product, nor yet do they say the opposite (as actually happens). Nor do the theories seem for example to account for, or to contradict, or in any way to relate to, a finding such as that buyers of the *small* pack-size of a brand buy it on average about as often as buyers of the *large* pack-size buy that (cf. § 10.3). And so on. In general then, these theories of consumer behaviour do not seem to predict any of the quantitative findings described in this book, or even their qualitative nature. They do not even state whether there should be a single, general pattern at all. Nor yet do they seem readily capable of suggesting *new* hypotheses to test, in areas which have not as yet been explored empirically.

For example, do the light buyers of a given brand buy so little of it because they are really heavy buyers of some *other* brand? Or are they simply light buyers of the total product-class (with *heavy* buyers of a brand also being heavy buyers of other brands)? Or is there some other simple pattern, or is there no general result at all, i.e. "it all depends on a thousand and one other variables" (the brand, the product, the time and the place, etc.)? At present, the empirical answer is not known, but the traditional "theories" do not tell either. A major effort of critical appraisal of such theories seems therefore to be needed – what *do* they in fact tell us?

A further line of academic model-building has been developed in terms mainly of buying variables as such, i.e. the buying and repeat-buying of a given brand or product and the switching to and from other brands. Included here are the kinds of "purchase sequence" models referred to briefly in § 1.4 of Chapter 1, Kuehn's often-quoted attempts at adapting certain forms of psychological learning theory to cope with buyer behaviour data [Kuehn 1958, 1962], and – probably best known of all – a variety of essays in applying simple Markov theory to brand-switching behaviour. Much of this kind of work has been summarised recently by Massy et al. [1970], and in review papers by Sheth [1967] and Lawrence [1966]. The work usually involves the *a priori* formulation of a mathematical model. This is then sometimes followed by an attempt to estimate the numerical parameters of the model with some actual data, which is simultaneously thought of as testing the model's goodness of fit or validity. Usually there is at most one such attempt at empirical analysis.

For the "learning" type of model initially explored by Kuehn, it is supposed that a purchase of Brand X should increase the consumer's probability of buying that brand again next time – an approach which is still popular [e.g. Aaker and Morgan 1971]. But Frank [1962] showed long ago that the mere existence of population heterogeneity – that some consumers may generally buy Brand X more frequently than others – would lead to an appearance of "learning" (thus echoing an old controversy over "proneness" and "contagion" in the study of accident statistics – see § 7.2). All the work in the present book seems to show that under stationary conditions, people whose last purchase was Brand X were more *frequent* buyers of X in the past and will therefore also be more frequent buyers of X in the future. No "learning" or systematic change in purchase probabilities needs to take place to account for the observed results.

The simple first-order Markov model – first put forward in the brand-switching context in the late 1950's [e.g. Lipstein 1959, and Maffei 1960] – is another example of an *a priori* approach to models of brand-choice. It specifies that the probability of switching between different brands depends on the specific brands as such, i.e. the switching probability from Brand A and to Brand B is so much, that from A to C so much, that from B to C so much, the repeat-buying probability for A so much, and so on, all irrespective of market-shares, penetration levels or purchase frequencies. And these switching-probabilities (p_{AB}, p_{AC}, p_{BC}, p_{AA}, etc.) are assumed to remain the same over time *. All this has in fact been assumed without benefit of empirical knowledge, essentially on the grounds that because buyer behaviour *might* behave like that, it would be useful to assume that it does (especially since the mathematics of simple Markov processes had already been learned at college). The Markov formulation has however little *a priori* analytic plausibility in the present context, as has been discussed more fully elsewhere [Ehrenberg 1965c, 1968f, 1971a, Massy and Morrison 1968, Charnes et al. 1971], and the *facts* of buyer behaviour as described earlier in this book now show that something like the exact opposite to the Markovian assumption is true: the incidence of repeat-buying and brand-switching generally does *not* depend on the specific brand as such, but on its current penetration level and average purchase frequency.

* These parameters will of course remain the same in the stationary situation, but so will any other aggregate statistic. The essential Markovian assumption is that the values of p_{AB}, p_{AC}, etc. remain the same over time even when market-shares (or the probabilities of purchase of each brand in each time-period) change with time.

11.3. Facts Without Models

The failure of past academic theories of buyer behaviour to connect successfully with the facts has not been due to any lack of facts as such. Indeed, many tens of millions are spent each year in commercial market research on collecting factual information about buyer behaviour, consumer attitudes, usage habits, exposure to advertising media, and so on – using consumer panels, sample surveys, product-testing procedures, etc.

This work is carried out for manufacturing firms and the like, for practical marketing purposes. The prime emphasis is generally on the quick tabulation and reporting of the latest facts. Little or no explicit effort is made to integrate different findings, or to learn from past as well as merely from the present experience. The orientation is often supposedly "problem-orientated". But solutions are sought on an *ad hoc* basis. It is ignored that the more *common* marketing problems are, by definition, recurrent (e.g. "How do I increase sales?" or "What is the effect of price/advertising/distribution?"). It is also ignored that for an answer to be valid in *one* instance, it must apply also in other instances (or that we need to know the reason why) and that only answers which are in fact generalisable to some known degree can possibly be of real use.

Whilst much information about buying and usage behaviour is collected in *ad hoc* sample surveys, the richest source of data on buying behaviour are the consumer panels in which buying information is collected continuously from the same sample of informants. The consumer-panel type of data are however used mainly for the analysis of aggregate *sales* (such as monitoring the sales volumes and the market-shares of different items, month by month and in the different sales areas of the country), and also to some extent for mapping out simple consumer profiles (e.g. the percentage breakdown of the buyers of each brand by age, size of household, socio-economic classification, etc.). These analyses make virtually no explicit use of the continuing "panel" nature of the data.

Tabulations concerning individual consumers' repeat-buying and brand-switching are relatively rare and isolated, although interest is growing *. Most such tabulations are of an *ad hoc* type. The most

* In the related field of measuring television viewing by panel-type operations, it was calculated some years ago that something like 500 million panel-type measurements a year were made in the U.K., of which 499½ million were not used for panel-type analyses [Ehrenberg and

concerted efforts probably stem directly or indirectly from the Kuehn-Rohloff type of "Gains-Loss" procedures [e.g. Kuehn and Rohloff 1965], where the results for individual consumers are aggregated after certain weightings (e.g. by the individual's own "market-shares"). This is done on the assumption that it will show what volume of sales one brand has in some sense "gained" or "lost" from each of the other brands. But the quantitative results produced depend on the assumptions and procedures adopted †. No explicit norms or regularities have actually been reported for interpreting what the tabulated results might mean. A recent version of broadly this type of approach has taken the major step of *explicitly* introducing certain "norms" by which to interpret the tabulated results [Buck and Jephcott 1971]. But the suggested "norms" are themselves only based on assumptions and have the property that they can virtually never fit the data [Goodhardt 1971].

One other commercial attempt to utilize consumer panel data in a "model-building" fashion is based on work by Fourt and Woodlock [1960] and Baum and Dennis [1961], developed further by Parfitt and Collins [1968], to predict the ultimate market-share of a newly-launched brand from early data. The evidence on the success of this type of procedure is apparently uncertain – it is quite widely used but there are cases where it does not work, and it is not known how to predict this. This is perhaps hardly surprising because the situation tackled is essentially one of a non-self-fulfilling prophecy, in that if the predicted sales-level is poor, something is done to try and change it, such as dropping the brand or increasing the advertising. But unpredictable failures to predict correctly tend to be particularly important with procedures whose aim is essentially predictive rather than diagnostic. In any case, it is typical of the ambitious aim of such more immediately "problem-oriented" work that the aim is in fact to forecast the penetration-growth of a new brand under the highly dynamic, non-stationary conditions of a new-brand launch (with a great many variable marketing inputs varying), whilst not yet having achieved (or utilising) any understanding or predictive ability concerning penetration-growth under the much simpler conditions of near-stationarity for *established* brands.

Twyman 1966]. The utilization of the panel data from the even more extensive "retail audit" types of operations that measure the flow of goods through retailers is almost certainly even lower.

† E.g. As sumarised by Lawrence [1966]: "Repeat-business for a brand is equal to the minimum of the share-points for the brand in two periods; an increase in share-points for a brand is allocated to the brand losing share-points according to the relative magnitude of their respective losses". What happens if one adopts a different allocation procedure?

11.4. From Facts to Theory, and Back Again

The approach described in the earlier part of this book differs both from the speculative kinds of academic theorizing and from the empirical extremes of the commercial data-collection procedures which have been briefly described in the two previous sections. The medium-term aim has been to develop generalisable knowledge and understanding of the facts, i.e. valid theory. However, the *start* was simply with the kinds of data — especially buying records from consumer panels — which were already being collected on a large scale in ordinary market research. The data clearly seemed "relevant" in some general way, but the initial work was undertaken with few if any specific pre-conceived hypotheses. The immediate aim was to look for any low-level patterns which held under a sufficiently wide range of conditions to be worth studying further. The nature of the available data meant that any apparent regularity could (at least in principle) be tested against innumerable sets of other data.

As one example of this low-level way of starting by just "looking for patterns" with a bare minimum of "hypotheses", we note that the origin of the work on the NBD and on its repeat-buying implications was virtually accidental. It arose from a practical problem-solving exercise, namely checking on a possible excess of "heavy" buyers in a certain product-field (as already mentioned in § § 4.2 and 6.4). It was known that the data in question were in some respects erroneous, in that the aggregate purchasing level of each brand was markedly too high when compared with outside check figures. It was suggested that if a theoretical distribution could be found which would fit the observed frequency distribution of purchases when the heaviest buyers were *excluded* (but would fail to fit when they were *included*), this would help to pinpoint that these particular data-errors lay in the "tail" of heavy buyers.

The finding was however that a particular theoretical distribution — the NBD — gave a good fit not only to the curtailed distributions (i.e. excluding the heaviest buyers), but also to the whole of the data (i.e. *including* the heavy buyers). This occurred for each of the brands in the product-field in question.

Once this regularity for the problematical product-field had been stumbled across, the question or "hypothesis" arose — turning the initial problem upside-down — whether the same Negative Binomial Distribution also gave a good fit to the observed frequencies of purchase in other, more "normal" product-fields *. The finding was that the distribution did in fact fit under a wide range of conditions, for the various specific pack-sizes of different brands in different product-fields, different length time-periods, different regions and countries, at different points in time, and so on, and that it has continued to fit (see Table 4.2 in Chapter 4).

Thus once an initial regularity has been found, it can start to act as a stimulus to searching firstly for further empirical patterns and then for interrelationships and for explanations. Some working hypotheses are therefore generated from the initial empirical findings. This *empirically* stimulated hypothesising however differs radically from the more purely *speculative* kind of theorizing – based mainly on the particular investigator's subjective insights – which has been referred to in § § 11.1 and 11.2.

A "deeper" question which arose from the empirical finding that the NBD gave a good fit to the observed frequency distributions was for example why all the data tended to follow a common pattern, and in particular why an NBD. A possible answer came from noting that there was a mathematical theory of an underlying stochastic process – of the Poisson-Gamma kind, as discussed in § § 4.5 and 7.2 – which might be relevant to consumer purchasing behaviour. The mathematics of the Poisson-Gamma model in question had already been formulated many decades before [Greenwood and Yule 1920], but this had been done in the quite different context of accident statistics, and essentially in the abstract, as one possible probabilistic or "stochastic" way in which an NBD *might* be generated, and with no reference whatever to consumers' purchasing behaviour. Recourse to such a prior *mathematical* formulation differs from the *a priori* theories of buyer behaviour referred to in § 11.2. Thus in the present instance there was already something for the mathematics to explain – a clearcut and generalisable empirical pattern in buyer behaviour which had to be accounted for, and which the particular stochastic model was in fact known to lead to, namely the highly specific finding that the NBD generally gave a good fit to frequency distributions of purchases observed in any single time-period.

The crucial part of the proposed mathematical model was the Poisson hypothesis. This was more advanced than the NBD when the latter is seen purely as a frequency distribution of the number of purchases made in a time-period. It essentially sought to show how the people

* It should perhaps be stressed that all these early findings were restricted simply to the frequency distributions of purchase as such, i.e. merely to the question of how many consumers bought an item 0, 1, 2, 3, etc. times in a time-period.

who made more than one purchase distributed these repetitive purchases over time. This suggested examining buying patterns over *different* time-periods, by way of checking whether theoretical deductions from the model actually held in practice. The first steps were of a relatively narrow or even perhaps arid statistical kind (such as the sampling error formulae illustrated by Fig. 6.1 in Chapter 6). But being successful, these results were followed within a few years by the first results on penetration growth (Table 4.11) and period-to-period repeat-buying (Table 4.7).

The mathematical testing and exploration of the Poisson-Gamma model in fact led to many new empirical discoveries, with "theory" now often suggesting hypotheses, i.e. new ways of looking at the (available) facts. A simple illustration is provided by the finding that the average frequency of purchase per "new" or "lapsed" buyer tends to be constant at about 1.4 (see § § 4.6, 7.6 and 8.4). This finding occurred several years after the NBD model had initially been put forward in the late '50's and when the LSD version of the theory had also already been developed [Chatfield et al. 1966]. The LSD theory then showed that the average frequency of purchase per "new" buyer was given by a formula which was rather simple when expressed in terms of the single LSD parameter q, namely $q/\ln(1+q)$. It was next noticed, by straight numerical analysis, that this theoretical quantity $q/\ln(1+q)$ varies very little for different values of q – namely between about 1.35 for $w>2$ and an absolute maximum of 1.43 – so that its value could generally be taken as being "constant" at about 1.4 (see § 8.6).

It was this sequence of *theoretical* findings concerning "new" buyers, and especially the simplicity of the final theoretical "constant", which then led to a re-examination of fairly extensive tabulations of empirical data made previously. Here the average frequency of purchase of "new" or of lapsed buyers had already been recorded (as one of a variety of figures) without their however ever having been explicitly "analysed" so as to see what pattern the figures might contain. Given now that the *theory* suggested a very simple hypothesis – i.e. that the figures should be virtually constant at 1.4 – it was noticed that the empirical values already tabulated were in fact generally about 1.5, i.e. close to the theoretically predicted value [Chatfield et al. 1966]. In subsequent work, the observed values again tended to be close to the theoretical value (as is illustrated in Tables 2.1 and 5.1, a systematic but rare exception being illustrated in Table 3.8 in Chapter 3).

The "1.4" result might of course have been noticed by direct empirical analysis in the first place (just as many other regular patterns have in fact been). Typically however, even the simplest pattern (like something being constant) is often only noted because of some prior suggestion – either theoretical or observational – that there might exist an answer which is simple and regular enough to be noticeable if one were to look for it. The "just noting something" does however remain the crucial step at this level of work. Thus the LSD version of the repeat-buying theory itself also arose accidentally, as an unexpected *positive* by-product of looking for a particular answer to one of the earliest *discrepancy* problems which occurs in fitting the NBD itself. The so-called "variance discrepancy" (§ 7.8) – i.e. the failure of the negative binomial distribution to fit in a certain respect – had led to the "hypothesis" (no more than an "idea", really) that the number of zeros in the observed frequency might be wrong (a matter of defining the population of consumers). It was quickly found that this hypothesis could not in fact explain the variance discrepancy at all. It was noted at the same time that much purchasing data was of a kind where the number of zeros did not greatly affect the fit of the NBD anyway, and that the data often fell within a certain range of parameter values where the LSD would act as a close (and much simpler) approximation to the NBD.

The major consequence of this quasi-accidental discovery that the LSD was appropriate was to show that one could analyse the proportion of buyers (b) quite separately from how often the buyers bought. This focussed attention on the average frequency or rate of buying *per buyer* (w), which in turn was the key to developing a great variety of further results, including all the *multi-brand* results summarised in Part V. The sequence of events here was briefly as follows. Initially, the study of repeat-buying had been restricted to specific pack-sizes of a brand, and was carried out not in terms of purchase occasions as nowadays, but in terms of the *number of units* bought (packs, bottles, etc.). This was slightly inconsistent with the basic formulation of the Poisson distribution, but it worked (and it related directly to *sales volume*, an immediate practical measure). In retrospect, it can be seen that this approach worked because for the products covered, only a single unit would be bought on most purchase occasions (cf. Table 3.3 of Chapter 3), although this very simple regularity was not known at the time.

In about 1965, Mr. L.J. Rothman fitted NBD distributions to petrol purchases and used *purchase occasions* (not "units bought") as the analysis variable [S.R.S. 1965]. This was for the simple reason that with petrol (where several gallons – i.e. "units" – tend to be bought per purchase), only something like a purchase occasion approach could – and did – work. (Because of this, it had not been noticed that in the earlier publications on the NBD, the number-of-units-bought formulation had been used). The success of Rothman's result for petrol led to a critical reappraisal of the earlier work. It became clear that the purchase occasion approach could work *at least* as well in the NBD theory for individual pack-sizes as the units-bought one had. Furthermore, the purchase occasion approach for the first time made it straightforward to deal with the purchasing of brands made up of different pack-sizes (as has also been discussed independently by Grahn [1969]). Thus it was found that by treating the amount bought per purchase as a *separate* variable and therefore working in terms of *frequency of purchase*, purchases of a brand made up of different pack-sizes (or varieties, etc.) could also be dealt with by the NBD, as could the purchases of many product-classes made up of different brands. These were *empirical* results. They were suggested by prior knowledge or low-level theory as just indicated, and they led in turn to deeper questions of *why* such aggregation worked (something which began to be answered only much later – see § 11.5 below).

The wide range results in the study of *multi-brand* buying which were opened up by the "purchase occasion" approach were also in the first place essentially *empirical*: thus, working in terms of w, the average frequency of buying a brand per buyer of the brand, soon made it obvious that its numerical value did not vary greatly from brand to brand compared with the differences in market-share, sales levels or penetration figures (as was discussed in § 10.2). The distinction between the traditional way of reporting sales and penetration figures (in grossed-up terms of value or weight etc. and total number of buyers), and the simplification produced by the purchase-occasion approach is illustrated in Tables 11.1 [and is discussed in more general terms elsewhere, Ehrenberg 1975a].

The simplicity of finding that this particular rate of buying w hardly varies also suggested looking at *other* rates of buying. Perhaps there were simple patterns also for *pack-size* rates of buying, rates of buying the *product*, rates of buying *other* brands, rates of buying per *sole* buyer of a brand, and so on? In all cases, the almost immediate findings were again very simple regularities, as has been set out in Chapter 10. The question raised by this increasing array of simple empirical regularities then became increasingly one of how they themselves were inter-related and why they occurred, and this led to theorizing at a somewhat higher level.

Many of the separate findings so far have in fact been found to interrelate and to fit together. They are beginning to build up to a more

Table 11.1. One Way of Setting Out the Sales Structure of a Market: Annual Sterling Sales and the Absolute Numbers of Buyers

Brand and Pack-Size		Buyers in '000*	Sales in £'000
Brand A	15*p*	3,361	3,332
	10*p*	2,169	1,106
	5*p*	1,337	308
Brand B	10*p*	2,025	1,095
	5*p*	1,260	287
Brand C	12*p*	1,538	959
etc.			

* Total population: 7,000,000.

Table 11.1a. An Alternative Presentation: Per Adult Penetration and the Average Rate of Purchases Per Buyer

Brand and Pack-Size		% Adults Buying (b)	Av. Purchases per Buyer (w)*
Brand A	15*p*	48	7
	10*p*	31	5
	5*p*	19	5
Brand B	10*p*	29	5
	5*p*	18	4
Brand C	12*p*	22	5
etc.			

* Rounded to 1 significant figure, to emphasise the relatively low degree of variation in the w's about the average value of about 5.

understandable whole. In the remaining sections of this chapter we briefly discuss examples of this (and of the remaining gaps). Some of the details are necessarily mathematical and research-orientated, but the general drift of the development is of some wider importance.

11.5. Reasons Why

The examples just discussed illustrate how prior knowledge (or "theory") suggested new hypotheses and new things to look at. But it was all at rather a low level – many of the results in question were uncovered almost accidentally (almost *despite* some of the hypotheses posed, rather than because of them). In virtually no case did it really matter whether or not the "expected" result was in fact found to be so, but just that something simple and general was found at all.

We now discuss examples of more advanced forms of theoretical work. Here it *would* matter if a theoretical hypothesis fails to fit the facts, for the important reason that the hypothesis is itself already a direct reflection or deduction from other facts. *Failure* in such more advanced theorizing therefore means that there is a new empirical discrepancy which needs to be accounted for. *Successful* theorizing of this kind on the other hand leads to integration – it helps to simplify, by explaining different findings in terms of each other.

The major example of such an integrative step described in this book is of course represented by the multivariate NBD theory itself (Chapters 4 and 7), namely that all the different aspects of repeat-buying behaviour are interrelated and that they can all be modelled by one single theory.

The theory is very simple partly because of all the variables or concepts that do *not* enter into it. For example, the NBD model implies that under stationary conditions, no "learning" or other systematic change in people's behaviour has to be allowed for ("new" and "lapsed" buyers in any given time-period for example are simply people who *regularly* buy the item rather infrequently). More generally, there is very little scope for other variables to enter directly into the situation, since the NBD model has no numerical coefficients to be accounted for in terms of other variables.

The only variables which enter into the NBD/LSD repeat-buying model are *buyer behaviour* ones, and even then it is only ones relating to the brand in question, Brand X say. Thus one of the principal variables which does not enter into the model is whether or not each buyer of X also buys other brands, Y, Z etc. This simple finding had for many years to be accepted merely as a matter of observed fact: the patterns of repeat-buying behaviour of the given Brand X could in day-to-day practice be successfully predicted *without* taking the other brands into account *. A theoretical explanation for the finding arose only ten years after the event, when the results represented by the duplication law $b_{XY} = Db_X b_Y$ of § 10.5 were established. Thus it became apparent that the correlation between whether or not consumers bought Brand X

* The case which was described in §5.6 is a very rare and quite recent exception, occurring because an additional factor – a specific taste-preference or the like – was at work.

in the analysis-period and whether or not they bought Brand Y in the period was generally very low, usually at most of the order of .1 or .2 or so *. It follows that if a model for repeat-buying of Brand X *did* explicitly try to take the buying of other brands into account, the coefficients in the resultant equations would be near-zero anyway. The duplication of purchase law therefore "explained" this major aspect of the NBD theory.

The duplication of purchase law $b_{XY} = Db_X b_Y$ itself has an outstandingly simple feature. Thus it relates b_{XY}, the proportion of the population who buy both X and Y each at least once in the analysis-period to b_X and b_Y, the proportion who buy X at least once and the proportion who buy Y at least once – without having to take account of how *often* any of these buyers buy X or Y (or anything else). The explanation of this empirical finding is now also known. It is provided by such simple facts as that w_X, the average frequency of buying for Brand X per buyer of X, varies little from the value of the corresponding rate w_Y, i.e. $w_X \doteqdot w_Y$. And similarly, the average rates of buying by duplicated buyers of X and Y, and the average rates of buying by *sole* buyers of X, are also approximately constant across brands (§ § 10.6 and 10.7). In any model for b_{XY} which *did* try to allow for differences in any of these rates of buying from brand to brand, the appropriate coefficients would once again be virtually zero, just because the rates do not vary. The simple nature of the relationship $b_{XY} = Db_X b_Y$ in this respect is therefore explained.

The relative constancy of the average buying rates across different brands which has just been referred to – i.e. that $w_X \doteqdot w_Y$ – is itself one of the most fundamental findings in the multi-brand part of the studies so far. Thus it links up the repeat-buying theory for different brands (since repeat-buying depends primarily on the value of w). It also provides a major marketing constraint (i.e. one clearly cannot increase sales in any *simple* way by merely getting existing buyers of a brand to buy it much more often – this would give an abnormally high w-value).

The $w_X \doteqdot w_Y$ result is a special case of the more general relationship $w_X(1-b_X) \doteqdot w_Y(1-b_Y)$ discussed in § 10.2 (see also Table 3.2). This has arisen as one of the more advanced pieces of theoretical work so far,

* The product-moment correlation coefficient is equal to $(D-1)\sqrt{\{b_X b_Y/(1-b_X)(1-b_Y)\}}$ Thus for two brands with (for simplicity) equal penetrations of .2, i.e. $b_X = b_Y = .2$, and with $D = 1.4$ as in Table 9.9, the correlation is about .4/4, which is .1. (For very high penetrations the true correlation is less than this formula would indicate on using the general D-value, as the D-value is then itself too high – see § 10.5.)

which it is worth outlining here. The point is that the $w(1-b)$ = constant form of relationship can be deduced theoretically from some quite *different* findings. The development was as follows. First came some empirical regularities for purchasing rates such as w_X and $w_{P.X}$, i.e. the average rates of buying the brand and of buying the product, per buyer of Brand X *. In time-periods of up to six months, it was noted that each rate did not vary much from brand to brand [Ehrenberg and Goodhardt 1968c, Ehrenberg 1969a], i.e. that

$$w_X \doteqdot w_Y \quad \text{and} \quad w_{P.X} \doteqdot w_{P.Y} .$$

This led Dr. C. Chatfield to note that if one made the simplifying assumption that buying of X and buying of Y was independent, then the latter type of finding ($w_{P.X} \doteqdot w_{P.Y}$) could not be true if the former finding ($w_X \doteqdot w_Y$) was. By this time, empirical results for longer time-periods had however begun to show that w_X *did* in fact vary somewhat from brand to brand (even if only relatively little), whilst the product rate of buying $w_{P.X}$ continued to show up as rather constant. Mr. G.J. Goodhardt therefore inverted Chatfield's earlier argument, and obtained the relationship

$$w_X(1-b_X) = w_Y(1-b_Y) = w. \text{, a constant}\dagger.$$

(The detailed mathematics are quite simple and are set out below.) The theoretical analysis therefore showed that w_X should be related to b_X, and it also gave the functional *form* of the relationship involved (i.e. that w_X varied with b_X as $w./(1-b_X)$).

It was found that this $w(1-b)$ relationship gave a good fit to an increasingly wide range of data, and that it was essentially consistent with the apparently contradictory earlier result that $w_X \doteqdot w_Y$ (i.e. that w_X did *not* depend on b_X). Thus *if* b_X and b_Y, the proportions buying X and Y in the analysis-period, are numerically small (e.g. 0.2 or less), the expressions $(1-b)$ are close to unity and the law simplifies to the statement that the average frequency of buying X per buyer of X is approximately equal to the average frequency of purchase of Y per buyer of that, i.e. that

$$w_X \doteqdot w_Y .$$

* $w_{P.X}$ means the average rate of buying the product (*P*) per buyer, *given* that he bought Brand X. The average rate of buying Brand X itself could correspondingly be written as $w_{X.X}$ (but is generally simplified to w_X).

† The constant $w.$ may be taken to be the limiting value of the average rate of buying a brand, as the proportion of the population buying it tends to zero.

This is essentially what occurs in relatively short time-periods when the numerical penetrations b are numerically small. Tables 11.2 and 11.3 set out the numerical results for the four leading brands analysed in Chapter 3 (Brand E was noted there to be atypical). Clearly the w's in 4 weeks vary very little and those in the 48-week period vary more markedly (although by far less than the differences in market-share – a ratio of 2 to 1 against one of 10 to 1 for Brands A to D). However, the $w(1-b)$ formulation in Table 11.3 stabilizes *both* sets of figures.

The mathematical argument in deriving the $w(1-b)$ result from *theory* runs as follows:

We start with the observed fact that the average number of purchases of the *product* made per buyer of Brand X is (to a close degree of approximation) the same as the average number of purchases of the product made per buyer of Brand Y (see § 10.4).

Supposing that there are only three brands X, Y, and Z (the argument can readily be extended to more brands), then for a population of N consumers there are Nb_X buyers of Brand X in the period, and their purchases of the product-class (i.e. purchases of Brands X, Y and Z) are made up as follows:

– firstly, Nb_Xw_X purchases of X, given that the Nb_X buyers buy X on average w_X times.
– plus $Nb_{YX}w_{Y.X}$ purchases of Y, where Nb_{YX} denotes the number of buyers of X who have also bought Y, and $w_{Y.X}$ is the average number of times these "duplicated" buyers of X and Y actually buy Brand Y in the analysis-period.
– plus $Nb_{ZX}w_{Z.X}$ purchases of Z, in corresponding notation.

Since the product-rate of buying per buyer of X and Y are the same, we therefore have (cancelling first N throughout) that

$$\frac{b_Xw_X + b_{YX}w_{Y.X} + b_{ZX}w_{Z.X}}{b_X} = \frac{b_Yw_Y + b_{XY}w_{X.Y} + b_{ZY}w_{Z.Y}}{b_Y}.$$

Now if buying of Brand X is independent of buying Brand Y, we have $b_{YX} = b_Yb_X$ and $w_{Y.X} = w_Y$, (i.e. the proportion of the buyers of X who buy Y is as in the population as a whole, and so is their average frequency of buying Y), and so on for the other terms. We therefore can write

$$\frac{b_Xw_X + b_Yb_Xw_X + b_Zb_X\,w_Z}{b_X} = \frac{b_Yw_Y + b_Xb_Yw_X + b_Zb_Yw_Z}{b_Y}.$$

Cancelling the b_X and b_Y respectively on each side, and also eliminating the identical terms b_Zw_Z, we have

$$w_X + b_Yw_Y = w_Y + b_Xw_X\,,$$

or collecting terms in X and Y,

$$w_X(1-b_X) = w_Y(1-b_Y)\,,$$

the result already mentioned.

Table 11.2. The Proportion of Buyers and the Average Purchase Frequency per Buyer, for the 4 Leading Brands in a Product-Field

(4-, 12- and 48-weekly results from Tables 3.1a and 3.2a in Chapter 3)

Period in Weeks		100b: % of Buyers			w: Av. Frequency of Purchase		
		4	12	48	4	12	48
Brand A	(50%)*	28	42	62	1.8	3.7	10
Brand B	(13%)	9	17	32	1.5	2.5	5
Brand C	(7%)	4	5	17	1.7	2.8	5
Brand D	(5%)	3	6	14	1.5	2.5	4
Average		11	17	31	1.6	2.9	6

* Market-share.

Table 11.3. The Proportion of Non-Buyers and the "Adjusted" Average Purchase Frequencies

Period in Weeks	100(1−b): % of Non-Buyers			$w(1-b)$		
	4	12	48	4	12	48
Brand A	72	58	38	1.3	2.1	3.8
Brand B	91	83	68	1.4	2.1	3.4
Brand C	96	95	83	1.6	2.7	4.1
Brand D	97	94	86	1.5	2.4	3.4
Average	89	83	69	1.5	2.3	3.7

This relationship between w and b was therefore initially obtained from the fact that the average rate of buying the product-class does not vary from brand to brand, together with the assumption that buying of one brand and buying another brand is independent. The latter is however no longer altogether an assumption. Instead, it has more recently been supported by a good deal of empirical evidence (to a close degree of approximation). Firstly, the duplication law $b_{XY} = Db_Xb_Y$ implies a very low (near-zero) correlation between X and Y, as already noted. Secondly, it is generally found that $w_{X.Y}$ differs little (if at all) from w_X; thus in general we can write $w_{X.Y} \doteqdot Cw_X$, where C is a number which tends to lie between about .8 and 1.0 for different product-fields (see § 10.6). In other words, buying or not buying of Brand Y has little influence (in correlational terms) either on the proportion of consumers who buy Brand Y or on how often they buy it.

If we work directly with this empirical evidence, rather than with the stronger assumption of complete independence, the same algebraic arguments as shown above lead to equation $w_X(1-CDb_X) = w_Y(1-CDb_Y)$. Since the numbers C and D are generally close to 1 (C being usually a little less than 1, and D a little bigger than 1), this relationship differs little from that of the $w_X(1-b_X)$ form *.

* For example, for the product-field analysed in Chapter 9, C is about 1.0 and D is about 1.4, so that we have $w_X(1-1.4b_X) = w_Y(1-1.4b_Y)$. In many other cases (see § 10.6), C is more like .8 and with a value of D still at 1.4 one would have $w_X(1-1.1b_X) = w_Y(1-1.1b_Y)$.

The relationship is in any case symmetrical (i.e. b_X and b_Y are subjected to the same multiplier CD). As a consequence, the $(1-CDb_X)$ and the $(1-b_X)$ types of relationships will broadly speaking give an equally good fit to the kinds of observed data available (e.g. as in Tables 11.2 and 11.3). For practical purposes at this stage, one however uses the $w(1-b)$ formulation which is not only simpler but also more general (i.e. no numerical coefficient has to be estimated in the individual case).

The important conclusions from this kind of analysis are therefore two-fold. Firstly, it is possible by *theoretical* argument to predict something very much like the right answer for the empirical relationship between w and b. Secondly, not only could such a new result be successfully *predicted* (a once-and-for-always aspect of how the result was discovered in the first instance), but of more lasting importance is the fact that three quite different empirical findings about buyer behaviour are found to be inter-related, namely the way w and b are related, the constancy of the *product* rates of buying, and the approximate independence of buying one brand and another.

The *precise* form of the relationship between w and b is however by no means settled. For example, the $w(1-b)$ version depends on zero correlations (i.e. both C and D equal 1), which is not *quite* true. Again, given that $w_X \doteq w_Y$ (either to a coarse or often to a good degree of approximation – see for example Table 11.2), it follows that the numerical value of the parameter a of the NBD/LSD theory in Part IV will also be about the same for different brands. But if $a_X \doteq a_Y$, it has been shown by G.J. Goodhardt that we should have $w_X(1-\frac{1}{2}b_X) \doteq w_Y(1-\frac{1}{2}b_Y)$. This is not the same as the $w_X(1-b_X) = w_Y(1-b_Y)$ result but rather similar to it, and therefore also gives quite a reasonable fit to the empirical data*. It also provides a possible link with the NBD/LSD repeat-buying theory – which so far has not entered into the across-brands w and b relationship at all – and is therefore one of a number of possible avenues for further exploration into the detailed connections between repeat-buying and brand-choice theory.

The empirical multi-brand results just referred to – namely the lack of any sizeable correlation between buying X and Y, and the fact that generally $w_X \doteq w_Y$, and that hence $a_X \doteq a_Y$ to a close degree of approximation – also explain another finding in the repeat-buying studies. This is for the situation where the purchases of two (or more) brands each of which follows approximately an NBD are combined. It has been found

* At this stage no systematic attempt has been made to establish which of the different formulations of the relationship might fit all the available data "best". Small differences in goodness of fit are not what is at issue.

empirically that the aggregate will generally also follow an (approximate) NBD. The necessary and sufficient *theoretical* conditions for the precise aggregation of NBD's were twofold, namely that buying of the two brands be independent, and that $a_X = a_Y$ (see § 7.7). These two conditions are now well supported by the facts, so that the aggregation of NBD's links up rather closely with the multi-brand findings already established. Similarly, the fact that brands made up of pack-sizes follow (approximately) an NBD when the individual pack-sizes are NBD, is similarly accounted for by the remarkable similarity of the average pack-size rates of buying across pack-sizes (§ 10.3).

And so one can go on: for example, the fact (§ 10.7) that the incidence of sole buyers varies across brands as $b_X/(1-b_X)$ follows from the virtual lack of between-brand correlation *, whilst the fact that the sole buyers' average rate of buying does not enter into this result is explained by the relative invariance of this buying rate from brand to brand (Tables 9.5 and § 10.7).

Increasingly therefore, the various initially *isolated* empirical results are beginning to inter-relate and to help explain each other. This leads to a theory which is parsimonious in terms of theoretical constructs or explanatory variables, which contains few if any unexplained numerical coefficients, and in which mathematics generally enter only at the later, more advanced, stages. (In contrast, the initial empirical patterns are obtained without any special analysis techniques, it being merely a question of finding simple regularities, typically that some rate of buying is constant – or *more or less* constant – across brands or pack-sizes or whatever **.)

More generally, the sequence of results discussed here appears to represent the three stages through which scientific theory usually passes. It would be difficult to improve on the way Sir Cyril Hinshelwood –

* The proportion of the population who only buy Brand X in the analysis-period is equal to b_X, the proportion buying X at all, times (on the indepence assumption) the proportions of the population who do *not* buy Brand Y, Brand Z, etc. i.e. $= b_X(1-b_Y)(1-b_Z)\ldots$. This equals $\{b_X/(1-b_X)\}\{(1-b_X)(1-b_Y)(1-b_Z)\ldots\}$ where the right-hand term is the same for all brands, i.e. constant, so that the proportion of the population who are sole buyers varies as $b_X/(1-b_X)$. The "constant" term here equals the proportion of the population who in the analysis-period do not buy the *product* (i.e. any brand), and it is instructive to compare the "Any Brand" line in Table 3.1a with the results in Table 9.4.

** This lack of statistical search procedures contrasts with the modern popularity of using the complex statistical techniques of multiple regression, factor and component analysis, cluster analysis and numerical taxonomy, etc. to try and find patterns in data at the *earliest* stages of any new study. But although it has been tried out on a very large scale over the last few decades, this "statistical" approach has shown neglible results. [Fuller discussions of some of the methodological issues are given elsewhere, e.g. Ehrenberg 1963, 1968c, 1969a,b, 1973.]

who was uniquely President of the Classical Association and of the Royal Society in the same year – put it [Hinshelwood 1957]. The first stage of scientific theory he described as

> "gross over-simplification, reflecting partly the need for practical views and even more a too enthusiastic aspiration for the elegance of form".

This is amply fulfilled by the empirical laws – such as the "constancy" of the various rates of buying – which have been described here. In the second stage of a scientific theory

> "the symmetry of the hypothetical system is distorted and the neatness marred as recalcitrant facts increasingly rebel against uniformity".

This is typically shown by the various departures and discrepancies from the simple "constant" relationships. In the third stage

> "if and when this is obtained, a new order emerges, more intricately contrived, less obvious and with its parts more subtly interwoven, since it is of nature's and not of man's conception".

Even here, aspiration for simplicity and some elegance of form – say that $w_X(1-b_X) = w.$, a *constant* – still tends to recur.

Over-simplification at the first stage does not mean that the empirical laws in question are not true – the purchasing rates in question *are* constant, within the limits of approximation that have been stated. The essential over-simplification resides in the concepts in which such laws are expressed. Hinshelwood's point is that any such law is unlikely to mean what it initially seems to mean. He quoted Alice: "Somehow it seems to fill my head with ideas – only I don't know what they are". This one increasingly learns as more and more recalcitrant facts accumulate.

11.6. Discrepancies

In looking to future developments, one problem area lies with the discrepancies or failures in the existing theory. (Two other questions are of filling *gaps* in the theory as discussed in § 11.6, and developing *applications* as discussed in Chapter 12.)

In repeat-buying itself, there are three discrepancies of a very systematic kind. Right from the beginning there has been the so-called variance discrepancy, namely that when the frequency distribution of purchases has a large variance, the theoretical value is even larger than the observed one (e.g. Fig. 8.1). For the kind of product where there is something of a weekly purchasing cycle, this discrepancy is now known to be due to the fact that few people buy more often than once a week. There is then a short-fall of people buying more than 12 times in say a 12-week period (compared with the theoretical NBD). This discrepancy only makes itself manifest when there is in fact an appreciable proportion of people buying the item virtually every week (§ 7.8).

A second well-established discrepancy is that in very short time-periods, purchasing behaviour is quite different from that in longer time-periods. Taking as an example products where there tends to be a weekly purchasing cycle, people who buy twice in two weeks will almost all buy once in each of the two weeks, rather than either buying twice in one week and not at all in the other, or buying twice the normal amount on a single purchasing trip. (This is quite different from the observed pattern in *longer* time-periods, where of the people who buy twice in two *months,* say, only about 50% will buy once in each month, and 25% will buy twice in the first month and 25% twice in the second month, in line with the independent-random-purchase-events model exemplified by the Poisson distribution). The week-by-week pattern which actually occurs therefore leads to a low average purchasing frequency per buyer in any single week (e.g. a w of 1.1 in Table 3.2a) and a rather boosted incidence of repeat-buyers (e.g. an average of 33% in Table 3.6a), whereas the incidence of repeat-buying predicted from the low w itself is very low (virtually zero) *. It is therefore a case of buyer behaviour being quite different in very short time-periods, and the NBD/LSD repeat-buying theory for time-periods of medium or longer length does not apply **.

A third discrepancy is that the NBD theory does not adequately describe buyer behaviour for certain items which at this stage may perhaps best be described as "very frequently-bought". Not enough is yet known to categorize these cases more precisely, but examples are the buying of the total product as such for certain product-fields, and the buying of an individual brand in other product-fields such as cigarettes, bread, etc. Here repeat-buying tends to be more regular than the NBD predicts.

* The NBD theory implies that if nobody buys more often than once in a given time-period, so that $w = 1$, then none of the buyers will buy again in the *next* equal period.

** It is obvious that neither the normal longer-term repeat-buying patterns nor the corresponding theory can apply in really short time-periods (e.g. of an hour-by-hour basis). This can also occur in longer time-periods for products such as items of clothing, where there is a minimum inter-purchase cycle which is far longer than a week (see §5.3 in Chapter 5).

In the last two cases it is not so much a question of the theory "breaking down", but rather that certain kinds of buyer behaviour (in very short time-periods or for very frequently-bought items) differ from the patterns of buyer behaviour which are more generally observed and which the NBD theory describes successfully. The existing theory helps to pinpoint the difference.

A new kind of theory is therefore required to model not only all these "discrepant" kinds of buyer behaviour but also all the repeat-buying phenomena which the NBD/LSD theory already copes with. One hopeful sign is that the three discrepancies outlined above all are themselves related, and appear therefore to be part of the same kind of general problem. Thus there is a shortfall of very frequent purchasers, and this is because for many product-fields, few people buy more than once a week – whereas for products where people *do* buy very frequently, buying tends to be more "regular" than normal. It would therefore appear that one single new theory should cope adequately with all these different discrepancy problems.

For many practical purposes, these discrepancies however do not greatly matter. This is particularly so because they tend to be at the "boundary" of the existing area of knowledge, i.e. for very heavy buyers, for very short time-periods, and for very regularly bought products. (It is rather like the Newtonian theory of gravity breaking down for very high speeds, and for subatomic distances, i.e. also at the boundaries.) It would be helpful to have a better theory, but this does not appear to be a matter of the greatest practical urgency: filling the *gaps* in the existing knowledge and developing its practical applications further would seem at this stage to be more important.

Two examples of work in the discrepancy areas, one negative and one with considerable positive promise, are however worth briefly mentioning. Firstly, a basic doubt about the NBD model is the Poisson assumption of §§4.5 and 7.2. Can one really assume that a given consumer's sequence of purchase over time occurs as if they were independent random events with the same fixed probability?

A typical query of the Poisson assumption occurs in W.J. Corlett's remarks in the discussion of the main paper on the NBD theory [Chatfield et al. 1966]. In relatively long periods of time (and given that the buying behaviour is stationary), the independent-random-events assumption is less unlikely than might at first sight seem, in that what might be thought of as positive serial correlation may largely be subsumed by one consumer's average rate of purchasing differing from another's. However, for very short time-periods the Poisson assumption certainly cannot fit (e.g. people are altogether less likely to buy in the middle of the night than during the daytime). But except for analyses of buying in very short time-periods as such, this need not have any adverse effects on the general (if approximate) fit of the Poisson-Gamma model. More recently, Herniter [1969] has also suggested the use of an Erlang distribution instead of a Poisson. The Erlang may be regarded as a Poisson distribution with every other event censored out and this introduces an element of negative correlation between successive purchases, i.e. one purchase inhibiting another for a certain period of time. This is promising but further work [e.g. Chatfield and Goodhardt 1972] and more general considerations of the nature of the observed discrepancies implies that far more than this would be needed to eliminate most of the discrepancy problems (see also § 7.8).

A second possibility for dealing with these discrepancy problems is to work in terms of purchasing *periods* instead of purchasing *occasions* (as has already been mentioned in the Foreword). Thus in analysing household purchases in a 12-week period, say, we would consider in how many *weeks* (from 0 to 12) the item in question was bought, rather than on how many occasions. This leads naturally to a distribution with an upper limit (here 12 for 12 weeks) and in this respect resembles more closely the kind of purchasing behaviour that is observed, i.e. very few people buying more often than 12 times in 12 weeks, with essentially everyone buying *at most* once per week *.

Some work has suggested that a Beta-binomial type of model will be useful here [Chatfield and Goodhardt 1970, and also Pyatt 1969 for an earlier but less specific approach using the Beta-binomial]. However, whether this particular mathematical model will cope sufficiently well with all or most of the outstanding problems requires further detailed study. If successful, this step would be like the earlier reformulation of the NBD repeat-buying theory from its initial formulation in terms of "numbers of units bought" to "purchase occasions" (§ § 1.4 and 11.4). The reformulated theory gave virtually the same results as the old one did in all those cases where the latter had already worked, but it also dealt successfully with discrepancies, and it allowed additional phenomena to be covered.

* One would allow for the *occasional* consumer who makes more than one purchase per week by an additional factor of the "average number of purchases per week", just as in the current work one can reconstruct sales volume by multiplying the average number of purchase occasions by a factor of the "average number of units bought per purchase" – see for example Table 3.3 and § 1.4.

Apart from the three very general discrepancies discussed above, more specific ones also exist. These are not many and – by definition – they are of course not general. Just because of their isolation, there is at this stage little chance of developing general theoretical explanations for them. (In any practical decision-taking situation, one has to do the best one can on an *ad hoc* judgment basis.) Two examples are

(1) The high frequency of purchase per "new" buyer which was found for all brands in the product-field analysed in Chapter 3 (Table 3.8).

(2) The *decreasing* incidence of repeat-buyers with increasing lengths of time-periods, in the pilot-scale data illustrated in Table 4.14 at the end of Chapter 4.

In both cases it is not so much a question of the observed data differing from the theory, but rather that the data in question differ from virtually all other kinds of observed repeat-buying (as generally successfully modelled by the theory). It will probably be difficult to find satisfactory explanations until *some* degree of generality for these occurrences has been established. But establishing under what particular conditions any such discrepancy tends to occur will of course begin to explain it – for example, the abnormally high frequency of purchase per new buyers reported in Table 3.7 has now also been observed in two other product-fields, and these are all *food* products. On a sample of three, there is therefore the suggestion that this phenomenon (where some people buy in unusually heavy "bursts" or "jags") is something to do with personal taste preferences and the like (the cloying of a flavour perhaps) rather than with external marketing pressures.

Discrepancies also occur in the multi-brand type of results outlined in Part V, as well as in linking the repeat-buying to the multi-brand results. In the latter context, the NBD results for the growth of the penetration b_T of a brand in time-periods of different length T and for the increase in the average frequency of purchase per buyer w_T (see Tables 4.8 and 4.10 or § 7.5) are for example not fully consistent with the between-brands formula $w_T(1-b_T) =$ constant. Further work in developing a *consistent* theory here is required.

A more immediate discrepancy in the multi-brand results themselves is that the duplication of purchase law $b_{X.Y} \doteqdot Db_X$ cannot hold for any brand X which has a high penetration b_X. For example, if $D = 2$ and $b_X = .6$ say (i.e. 60% of the population buying X in the analysis-period), then the proportion of buyers of some Brand Y who buy Brand X would be greater than 100% according to the law. In practice, the law in fact overstates the observed duplication levels for high-penetration brands (see for example Brand A in Tables 9.7 and 9.8) and the theory is clearly wrong. A better mathematical formulation is required, but no successful answer has yet been found. In looking for one, one may note that the duplication law states how many buyers of one brand also buy another brand whilst completely ignoring how *often* they do. This is explained (as already mentioned in § 11.4) by the fact that the average frequency of purchase by duplicated buyers of Brand X usually does not depend on which other brand one is considering. In effect, one is here dealing with the fact that $w_{X.Y} \doteqdot w_{Y.X}$, to a fair degree of approximation. However, for brands with a high penetration, this constancy of the average rates of purchasing of course no longer holds, as we then have the more general relationship $w_{X.Y}(1-b_Y) = w_{Y.X}(1-b_Y)$ operating (see § 10.6). This suggests that the reformulation of the duplication of purchase law which is needed to deal with high values of b_X will probably have to take into account not only the numbers of buyers but also their frequency of buying (for brands with such a high penetration level) *. Put in another way, the various empirical inconsistencies and sub-patterns appear to be coalescing and only a new mathematical formulation (a new "model") is needed.

11.7. Gaps and Future Developments

Although there are relatively few *discrepancies* in the existing results to be sorted out, many *gaps* remain to be filled and many are large ones, especially at the time of the first edition in 1972. Here we provide a brief listing rather than attempt any very full discussion.

Easily the most important gap in our existing knowledge relates to the penetration of each brand. Thus both market-shares and repeat-buying under stationary conditions turn on two main variables, namely the penetration b and the average buying frequency w of each brand. The buying frequency matters most in predicting repeat-buying, but does not greatly vary from brand to brand. In contrast, the penetration varies a lot from brand to brand (and hence primarily differentiates brands in terms of their market-shares), but does not greatly affect repeat-buying (as exemplified in its extreme form by the LSD version of the theory). And the two variables, b and w, are themselves inter-related by the $w(1-b)$ = constant type of relationship. So we are left really with one variable, the penetration b, from which w, and sales

* But one may be able to translate this back into terms of b, since $w_X = w./(1-b_X)$.

(wb), and all aspects of stationary repeat-buying can then be successfully predicted (in principle, and largely in practice). The big gap therefore is why one brand has more buyers than another. However, this is not so much a gap in stationary theory as a need for a different kind of theory altogether, namely one of consumer *dynamics* (how did one brand come to have more buyers than another?), and this is discussed in Chapter 12 as an *application* of the stationary theory.

A related question is why the average purchasing frequency of w of each brand is what it is. We know why it is more or less the same for different brands (from the theoretical deriviation of $w(1-b)$ in § 11.4).

Table 11.4. The Frequency Distribution of Purchase of the Large and the Small Pack-Size of a Typical Brand in a Half-Year

(Observed Values "O" and Theoretical NBD Predictions "T")

24 Weeks		Number of Purchases Made							Average per Buyer
		0	1	2	3	4	5	6+	
Buyers of									
The LARGE Pack-Size	O: %	81	9	3	2	1	1	3	3.0
(= 100%)	T: %	(81)*	8	3	2	1	1	3	(3.0)*
The SMALL Pack-Size	O: %	89	4	2	1	1	1	2	3.3
(= 100%)	T: %	(89)*	4	2	1	1	1	2	(3.3)*

* Used in fitting.

But in 1972 we did not know how to predict the value of w from *other* aspects of buying behaviour, such as the number of brands, the total rate of product usage, etc. – a gap which is now filled in Chapter 13. The $w_X \doteq w_Y$ type of relationship needs in any case further explanation, given that w is an *average* anyway (and one of a highly skew distribution), and given also such additional facts as that the average rate of buying the *product* is generally so very much higher (§10.4).

The same kind of question is raised, perhaps even more strikingly, by the fact that the average purchase frequency per buyer of the large and small pack-sizes of a brand tend also to be the same. Table 11.4 gives an example for a particular brand, where buyers of the large pack-size bought it on average 3 times in 24 weeks, and buyers of the small pack-size bought that on average 3.3 times (i.e. almost the same *average* rates). However, the table also shows how the distribution of purchase frequencies about these averages is very wide. The question is then why these *averages* are the same, especially when, on average, these buyers also buy *other* brands or pack-sizes so much more often?

The problems here are in the first instance amenable to straightforward analysis and exploration of the kind already outlined in this book. For example, it will be a matter of establishing by direct empirical analysis across a wide range of brands and product-fields in what ways buyers of one brand (or pack-size) buy *other* brands so much. Is it *light* buyers of Brand X who buy a lot of other brands, or is it *heavy* buyers of X who are also heavy users of other brands? Or what? At present, the answers are not known.

The questions raised so far are basic gaps in our knowledge of *multi-brand* buying, rather than gaps in our understanding of repeat-buying itself. But since they involve the basic parameter w, they do relate to the relationship between the repeat-buying levels of *different* brands. But even for the repeat-buying of any single brand, there is a good deal of scope for further expansion of detail. Examples are the repeat-buying behaviour of sub-groups of the population such as *sole* buyers [cf. Cannon et al. 1970 for an exploratory study here], repeat-buying by *heavy* buyers of the product-field and by *light* buyers of the product-field (as opposed to heavy or light buyers of the *brand*, which is already fully understood, e.g. Table 3.10 and § 7.6), repeat-buying by different demographic groups, and so on.

In the multi-brand or brand-choice area, there are other questions still to explore. Specific examples are the factors which determine the parameter D in the brand-duplication law $b_{XY} = Db_Xb_Y$, and the factors which determine (in a related sort of way) the incidence of sole buyers, and the levels of the various buying rates involved (e.g. $w_{X.Y}$, w_S, etc.). More generally, we have already noted in § 11.6 that the duplication law itself is not quite adequate (for high values of b_X). But apart from such a relatively minor quantitative problem, it might now seem that the law is itself very obvious – once one thinks about it, it is hardly surprising that the tendency of buyers of Brand Y also to buy Brand X is directly related to the *general* popularity b_X of Brand X in the population as a whole. However, this is no longer quite so "obvious" when seen against the fact that completely different kinds of results hold for certain aspects of *pack-size* duplication of purchase. Thus some initial exploratory work indicates that whereas duplication of purchase between different brands all of the same pack-size follows the usual duplication law, duplication of purchase of different pack-sizes of the *same* brand follow a completely different pattern, with the penetration of each pack-size (i.e. its general popularity in the population as a whole) having little or no role to play.

The relationship between pack-size loyalty and brand-loyalty – and similarly between variety (e.g. flavour) loyalty and brand-loyalty – are other related areas which still require a good deal of study and explanation. For example, to what extent (and under what circumstances) is there a higher tendency to switch to a different pack or flavour of the same brand rather than to the same pack (or flavour) of a different brand?

A further brand-choice area to explore is that of "switching" from one time-period to another, rather than that of duplication of purchase within a given time-period. The general pattern will probably be broadly the same as in the brand-duplication model, but the diagonal elements will certainly be different (100% in the case of within-period duplication, but less – i.e. the incidence of repeat-buying – in the case of period-to-period brand-switching).

Yet another area for exploration are any comparable patterns of repeat-buying and choice-behaviour in relation not to *brands* or the like, but to retail outlets (or *types* of outlets), together ultimately with the interactions between shop-choice and brand-choice (see also p. 249).

Underlying all this discussion of multi-brand buying behaviour is the need to develop a stochastic model for brand-choice formulated in terms of the individual consumer, comparable with and integrated with the Poisson-Gamma type of stochastic model on which the NBD theory of buying frequencies rests.

A completely different type of gap in the theories developed so far in this book are that they are all self-contained – relating different aspects of buying as such, with no *external* variables playing any role (e.g. attitudinal variables or marketing inputs). Furthermore, in the *repeat-buying* theory there are no numerical coefficients which could possibly be linked to outside variables *. However, it has already been established for frequently-bought products that buying (or "usage") behaviour is consistently related to such "external" variables as consumers' specific attitudes towards the brands, consumers' brand awareness, and their purchase intention. Thus it has been found that the percentage of

* In the multi-brand theory, there are a small number of coefficients, such as D in the duplication law, but it is almost certain that these can also be meaningfully linked with other buyer-behaviour variables as such (e.g. the number of brands available, their penetration levels, the average buying rate, etc.). Thus G.J. Goodhardt has noted the identities $D \equiv \{\Sigma b_{XY} - \Sigma b_X\}/\{\Sigma(b_X)^2 - \Sigma b_X^2\} \equiv \{\Sigma r_i^2 - (\Sigma r_i)\}/\{\Sigma(r_i)^2 - \Sigma b_X^2\}$, where r_i is the number of brands bought by the ith individual, but no effective study of this formulation has yet been made.

the population giving a particular attitudinal response to a brand (e.g. "is kind to the hands", or "is nourishing" or more generally, "intend to buy it") is directly related to the percentage of the population using the brand [e.g. Bird and Ehrenberg 1966a, b, 1970-2] †. It seems at this stage that knowledge of brand usage (or buying) will help to explain attitudinal responses, but ultimately one would perhaps hope for the direction to be in the usually more expected one of data on consumer's attitudes helping to explain their buying behaviour. It is presumably in terms of the individual consumer's attitudinal and other "non-buying" characteristics that one will have to seek to explain why it is that one consumer is a *light* buyer and another is a *heavy* buyer of a given brand, or why one consumer buys only one brand and another buys many brands. A more psychologically and sociologically (rather than behaviourally) oriented theory of buyer behaviour might perhaps be developed which can link up (and predict) many of the present findings, and therefore help to "explain" them‡. (See also Ehrenberg 1974, Barwise and Ehrenberg 1988 for discussions of the A-T-R theory of consumer behaviour).

11.8. Summary

In this chapter we have briefly discussed the background of *other* work into buyer behaviour, and the ways in which the findings reported here actually arose and how they have begun to interrelate.

Numerous speculative theories of buyer behaviour have been put forward in the past but it has not been clear how they relate to the facts, partly at least because the facts were themselves not organised. Now that a good many general patterns of buyer behaviour are known, it seems however that there is no relationship between the past theories and this empirically-based knowledge, nor yet do the theories seem immediately to suggest *new* hypotheses about buyer behaviour to explore.

† The data analysed usually involved a measure of brand "usage" rather than one of purchasing as such.

‡ A reasonable parallel here from physics might be that the pressure P times the volume V of a given body of gas is constant under fixed temperature conditions, i.e. PV = constant, a behavioural law which was "explained" only many years later by Avagadro's molecular theory of the behaviour of molecules in gas – but for practical purposes it is still usually the PV = constant type of relationship rather than *detailed* explanatory model which is of practical use.

The lack of proper connections between the past theories and factual knowledge has not been due to any absence of facts as such, since a great deal of data is collected in commercial market research. This factual information (rather than the past theories) has provided the starting-point for the work described here, which was undertaken as a search for empirical patterns without any formalised preconceptions of how or why people behave as they do. Typically, the starting-point of the work was a semi-accidental discovery, namely that the NBD gave a good fit. Empirical regularities of this kind then started to suggest new ideas and hypotheses for further study, often at an informal "let us look at so and so" level of suggestion. Only as more and more results accumulated did any formal theorizing and integrative work begin, the main example being of course that under a wide range of conditions, all the different aspects of repeat-buying itself can be subsumed by (or predicted from) a single theoretical formulation (i.e. the NBD model). Furthermore, later results in the area of multi-brand buying have started to explain many of the earlier findings.

A relatively small number of discrepancy problems remain in the existing theory of repeat-buying and in the multi-brand results, and these need to be followed up. They appear however at this stage to be less important than either the pursuit of previously unexplained aspects of buyer behaviour or the development of further practical applications of the results.

Added after proof in 1972:

Shop-Choice (see Paragraph 11.7)

In a recent paper ("Consumer Loyalty – A Fresh Look", read at the Annual Conference of the Market Research Society, Brighton, March 16th 1972), Mr. J.St.G. Jephcott has published some very promising results on shop-choice.

The initial findings (based on two product-fields over 24 weeks) show that shop-choice patterns for supermarkets are like those for brand-choice discussed in this book. Jephcott's comparisons include penetration growth (along the lines of Table 3.1a in Chapter 3), average purchase frequency (Table 3.2a), average frequency of purchase by repeat-buyers and "new" buyers (Tables 3.6 and 3.7) and the Incidence of 100%-loyal or "sole" buyers (Table 9.4 in Chapter 9).

If these kinds of results can be extended, it looks as though most of the present theory for brand-choice can be directly transferred to shop-choice also.

Added in 1987:

Store-Choice

In work on consumers' store-choice in recent years it has been found that the patterns are like those for brands. The same models hold. In brief, store-choice is like brand-choice.

Following Jephcott's start in 1972 as noted on the previous page, Wrigley (1980) and Wrigley and Dunn (1984a and b) of the Bristol Geography department have found that the repeat-buying and multi-store patterns predicted by the NBD and Dirichlet models apply at the level of individual stores. Kau and Ehrenberg (1984) showed that they applied also to national groupings of stores in the UK (chains etc), and Uncles and Ehrenberg (1987a) in the US.

At a further level of disaggregation, Kau (1981) has found that brand choice patterns within a given store chain also follow the NBD and Dirichlet.

CHAPTER 12

DEVELOPING PROTOTYPE APPLICATIONS

12.1. Problem-Solving and Basic Knowledge

A major division running through the work described in this book has been between gaining basic knowledge on the one hand, and trying to tackle practical problems on the other (the division between pure and applied science). The broad strategy has been to concentrate first on gaining some general knowledge and understanding of the system (here buyer behaviour), and only then to try and apply the knowledge to solving practical problems*.

In practice, the division is not always so neat – not *all* pure research precedes *all* practical applications. In carrying out basic research, practical problems provide pressure and also act as stimuli. Thus in the present instance, the whole work on repeat-buying started from a very narrow practical problem, namely trying to establish whether or not in one particular set of data there were "too many" heavy buyers (see § 11.4). This start from a practical problem worked out well in this instance, but this was a matter of luck. As a general strategy, it would be very risky to try and see whether there is something *abnormal* in some data (i.e. "too many" heavy buyers) without ever having looked to see what – if anything – might be *normal*. It would be safer to invest first in some basic research (what *are* the normal patterns?) before using the results for problem-solving. This approach would pay off in any subject-area where it is known that a variety of practical problems will in fact continue to arise.

The pay-off of such basic research will however only materialise at any reasonable speed if it is also recognised that there must be specific work on developing prototype applications of the new knowledge. One cannot expect a technologist, administrator, executive, manager, etc. (or even another scientist) to apply new basic knowledge if it has not been shown to what problem this knowledge can in fact be applied, and in what way.

* This strategy appears to run counter to the more popular modern "problem-solving" approach typified by much of the work in operational research for example, where practical problems are often tackled without as yet having any substantive prior knowledge of the system in question.

Investment is therefore required not only into basic research itself, but also into the development of prototype applications, so that usable results or procedures are already available when the need for more or less *routine* use in fact arises. The investment decisions here are complicated by the existence of sizeable time-lags: one must settle on "relevant" areas for study, then wait for results, then settle on potentially "important" areas for developing practical applications, and again wait for results. Furthermore, a balance must be struck between investing in the various research possibilities and the different types of development work. This largely becomes a matter of timing and of priorities. For example, on the research front it is often better to tackle new aspects of a problem area rather than to pursue discrepancy problems arising out of recent work (however "interesting" and helpful a solution would be). Thus in the work described here, the decision was made (after a time) to start exploring multi-brand buying behaviour rather than following up with any great urgency the various technical problems which remained in repeat-buying theory (such as the variance discrepancy – see § 11.6). Similarly, an *early* start must be made in allocating effort, time and money to trying to develop some applications. The danger is that, having made the decision to invest in research, the parallel decision to start on the probably even more costly development of practical applications is delayed, so that a "technological gap" arises.

In the remainder of this chapter, we briefly discuss three broad types of practical applications, namely applications

– to management problems, both in marketing and more generally, in § 12.2,
– to reaching some general understanding of brands and of consumers, in § 12.3,
– to further research, both under stationary and under "dynamic" conditions, in § 12.4.

It should perhaps be stressed straightaway that "further research" is itself regarded here as one form of "practical application" of the results, as long as the further research does actually utilize the previous results.

12.2. Management Applications

Management problems in marketing typically concern decisions about product-formulation, price, distribution, promotion and selling activities. There are related problems on the governmental side, such as stimulating production and new product development, in controlling product-proliferation, in ensuring competition, in controlling promotional activities, and more generally in dealing with a "consumer-orientated" society. In as far as such problems all involve the consumer, knowledge and understanding of buyer behaviour should be of help in dealing with them. But the detailed mechanics of this are not necessarily immediately obvious.

Sometimes a particular result of research into buyer behaviour can have more or less *instant* technological implications. For example, the fact that the average frequency of buying per buyer does not greatly vary from brand to brand implies that one cannot increase sales (at least not *radically*) by simply making existing buyers buy more (because they won't). It also immediately tells us things which we can predict about the ultimate behaviour of a new brand or other item which one may be planning to launch [Ehrenberg 1971b and 1987]. Such immediate implications of a new research result are however unusual. More often, a new research finding does not seem to answer *any* practical questions. This is however because no relevant questions have in fact been asked.

A major step in developing practical applications is therefore first of all to pose some decision alternatives, namely that one might take either action P or action Q. Alternatively, one must pose different *aims*, e.g. to try and increase sales by either attracting more buyers, or by making existing buyers buy more. Posing such alternatives then allows them to be evaluated against our knowledge of buyer behaviour.

Posing alternatives in terms of actions or aims tends in practice to be somewhat difficult, partly because one needs to know something about buyer behaviour in the first place to develop alternatives. For example, only after knowing that some buyers of Brand X are heavy buyers and others are light buyers of the brand could one formulate a policy of aiming one's advertising at increasing the number of heavy buyers. Next, by knowing that the distribution of light and heavy buyers of a brand tends to follow a highly predictable and set pattern (i.e. that of the NBD) one sees that one cannot aim simply to influence heavy buyers without having at least indirectly to influence the incidence of light buyers as well. Knowledge of buyer behaviour therefore does not provide an instant answer to the naive "how much should we spend on advertising" type of question, but it can show that a particular advertising aim is not worth pursuing. More generally, the existing knowledge imposes a great many constraints on possible action. The kinds of results described in this book therefore describe a part of the context in which marketing management has to take its decisions.

Table 12.1. Some Practical Problems tackled through Studies of Buyer Behaviour

New Brands
Test Marketing
Re-launches

The Growth of Penetration
The Leaky-Bucket Theory
Short-Term versus Long-Term Loyalty

The Effect of House-Names
Own Label Brands
The Level of Retail Distribution

Complementary versus Substitute Brands or Products
Product-Field Definition
The Profitable Length of the Product Line

The Spacing of Pack-Sizes
Price Differentials
Price Changes

The Way Sales Increase
Defending one's Brand-Share
Regional Sales Weaknesses

Impulse versus Habit-Buying
Household versus Individual Brand-Choice
Life-Cycle Assessments

Evaluating Individual Deals or Promotions
Below-the-Line Activity generally
Deal-Prone Consumers

Expertise in practical applications is something which has to be built up over time. Some cases have already been described in Part IV, such as evaluations of a new brand or a promotion or a seasonal trend. A variety of other applications have also been carried out, mostly as yet on an isolated "case-study" basis. Some of the more important topics are listed in Table 12.1. Much of this experience is as yet superficial and piecemeal, but then many of the research results (especially in the multi-brand buying area) are themselves a few years old. Few intensive attacks on particular problems have yet been made (the most developed type of application being the general mapping out of buyer behaviour in a given product-field, as illustrated in Chapters 3 and 9). In 1972 there had for example been no detailed investigation of the way in which the existing knowledge can help in guiding and evaluating advertising policy – a gap now closed (Ehrenberg 1974). But several studies of promotional activities and of product-proliferation are already beginning to add up to more systematic expertise in these particular areas. Two or three brief examples of current studies may help to clarify the issue.

Distribution Problems. A common marketing problem is that sales of one's brand are lower in one area of the country than in another. This is often linked to lower retail distribution (i.e. the percentage of shops stocking the item), leading to the chicken-and-egg problem of whether sales are low because of low distribution, or distribution is low because of low sales. This problem can be reformulated by asking whether sales are low because *supply* (distribution) is poor or because *demand* (buyer behaviour) is weak or "different". In this form, the problem becomes soluble: if (as found in a recent study) there is nothing in the repeat-buying and brand-switching patterns of one's brand which differentiates it from other brands in the area and from one's brand in its stronger sales area – *other* than that it has fewer buyers – then there is in that sense no *demand* weakness.

Another type of distributional decision-problem which arises is when a particular retail outlet (or chain or group of outlets) decides not to stock a certain advertised Brand X (or to give it less shelf-space), because a competitive manufacturer has just offered an attractive "bonus" to the retailer (i.e. a price-cut which may or may not be passed on to the consumer). The question then is who will actually lose by the retailer's decision not to stock Brand X for the time being. Will it be the retailer himself, because brand-loyalty is sufficiently strong for previous buyers of Brand X to go to another shop to buy it? Or will it be the manufacturer of Brand X because shop-loyalty is stronger than brand-loyalty, and consumers will switch to other brands in the given shop? How should the manufacturer of Brand X react to this situation or guard against it – by *also* price-cutting to the retailer, by strengthening the consumer appeal of his brand by increased media advertising, by promotions geared directly at the consumer, by just sitting tight, or what?

To start answering such a practical decision-problem, it helps to know something about shop-loyalty and its interaction with brand-loyalty as is now happening (see p. 241). We are beginning to know to what extent consumers normally buy the product at different shops, to what extent they buy the different brands at the same shop, how this varies for different types of products, for different types of shops, for brand-leaders and for less-heavily advertised brands, and so on. What are the various factors which determine shop-loyalty and its interplay with brand-choice?

What general patterns are there, and what exceptions? And what determines these patterns and these exceptions? The kinds of analysis used run parallel to those described in the present book for purchases of a given product-class when subdivided by its different brands, but now distinguishing purchases of the product made at different retail outlets. A second form of analysis is to relate such shop-loyalty results to those for brand-loyalty as discussed in this book. By knowing and understanding some of the ensuing results, one can clarify the manufacturer's or the retailer's management problems.

Length of Product-Line Decisions. Many manufacturers run different varieties or versions (such as flavours) of their brand. This raises decision-problems about the length and nature of the manufacturer's product-line. How many varieties should he run, especially of those selling less well? Should he aim to match the varieties offered by competitive brands, or should he fill the gaps? In launching a new variety, how much of this would simply cannibalise sales of his existing varieties? Can he afford *not* to launch new varieties? Should he launch another brand (repeating the most popular varieties) rather than add more varieties to his existing brand? Should he advertise his different varieties altogether under the general brand-name, or need he push them as individual items separately? Are there certain "pillar" varieties which everybody tends to buy regularly, and certain other items which consumers treat as occasional extras? Are there clusters of people who prefer particular groupings of items?

Here, analysis of repeat-buying and switching across both brands and varieties is needed. For example, do buyers of one's brand switch to other brands in order to buy varieties which the manufacturer does not make, or in contrast, do people switch to another brand mostly in order to buy the same variety again (as found in one recent study)?

Pricing. In many markets, there are more or less permanent price differentials between different brands (e.g. differences in quality, or advertised brands versus own-label brands) together with a great variety of temporary price-cuts. Repeat-buying analyses carried out at different price-levels may show that repeat-buying of *cut-price* items is low, as is that of full-price items, whereas repeat-buying of the total (irrespective of price) is normal. This would indicate that much of the switching is between prices within brands, i.e. that there is no price-segmentation. On the other hand one might find that switching is between brands within price, together with sub-groupings (e.g. in a four-price market there might be much substitution between the "expensive" and the "high-priced" items on the one hand and between the "medium" and the "low-priced" items on the other, but little between the two top prices and the two bottom ones).

In general, an initial prototype study of any particular problem area needs first to be made, to be followed by a more *systematic* range of studies concerning that problem, so as to provide real guidelines for practical decision-making. For example in §2.4 of Chapter 2 we illustrated the analysis of a particular seasonal trend for a particular brand (showing that that year there were all-the-year-round buyers of the brand who were unaffected by the trend, and peak-season-only buyers – See also Wellan and Ehrenberg 1987b). But the need now is to learn also what the seasonal patterns for that brand are like in *other* years, and for *other* brands in the same product-field, and for different product-fields altogether, and so on. Similarly, systematic studies of promotional activities, of new-brand launches, of own-label brands, of price differentials, of distributional effects, etc. are all needed.

In general, then, what is needed is to formulate problems and to get down to practical evaluative work against the existing knowledge of buyer behaviour. This involves a to-and-fro process of problem formulation, of provisional and partial answers, of further clarification of the problem and of further technical work. Hopefully, the process can often be made into a benign spiral rather than into a vicious circle where the researcher waits for the practical man, and the practical man waits for the researcher.

But one basic lesson is already clear. In dealing with routine marketing data on buyer behaviour, we no longer have to present vast arrays of more or less undigested figures, nor do we have to tackle all problems as if they were new. Instead, we can evaluate the data and the decision-alternatives against expected norms or predictions of what generally happens. This leads to the form of exception-reporting that has been illustrated in Chapters 3 and 9, and the kind of problem-solving illustrated in Chapters 5 and 6. Most of the figures are more or less as predicted and only the exceptional figures need *special* attention, whereas before – without interpretative norms – *all* the figures needed a judgmental type of interpretation. An important consequence of the relative lack of deviations from the norm – they tend to be few in number, are mostly fairly small, and often systematic if they occur at all – is to focus attention not so much on the irregularities as on the normal patterns of buyer behaviour. It is here that much more interpretation and more linking to marketing decisions is required.

12.3. The Brand and the Consumer

The second broad area of applications mentioned in § 12.1 was that of reaching some general understanding both of the consumer and of the role of competitive "brands" in consumer goods marketing. Any attempt to develop formal theories in these respects is beyond the intended scope of this book, but some brief remarks will show the possibilities of using the present results.

Thus one of the main lines of traditional thinking in marketing has been that the different brands in a given product-field have to possess different properties or attributes in order to appeal to their consumers. These differentiating properties may either be real, or be more of the "brand-image" type (i.e. attributes with which a brand is invested by its advertising, and general promotion, including packaging). This view – exemplified by the current popularity of the idea of searching for market segmentation, with different brands satisfying different "segments" of the population – has authoritatively been put as follows [King 1970, see also Jones 1986], in the context of the relationship between manufacturer and retailer:

> "In the current situation, the only leverage the manufacturer can apply to the retailer is his relationship with the consumer. And the main element in profit growth is going to have to lie in making his brand more valuable to the retailer through its being more valuable to the consumer. And that means his brand must be unique, it must have no direct adequate substitute – because it is in this that value lies after all. Sustained profit growth will only come if his brand has unique added values."

However, all the evidence indicates that as far as the consumer is concerned, different brands in general induce the same kind of buying behaviour, both in terms of repeat-buying and in terms of brand-switching. Furthermore, the same kinds of *attitudinal* responses tend to occur [Bird and Ehrenberg 1966a & b, 1970, Bird *et al.* 1969, Barwise and Ehrenberg 1985, 1987]. In general, the only substantive difference between one brand and another in terms of consumer response appears to be that one brand has more buyers than another (except where some *real* difference in product-formulation exists). In addition, some early experimental results suggest that the common forms of brand-loyalty as described here may perhaps be generated without special marketing efforts such as advertising [e.g. Ehrenberg and Charlton 1972]. There appears to be no evidence that one brand needs to differ from another in order to sell more.

As for the consumer, it is clear that his buying behaviour generally follows simple and predictable patterns. It seems to be characterised by regular habits, rather than by constant search or by uncertainty. The implication is that the consumer is less affected by the content of advertising and marketing action as such than is often claimed, but responds in a "reasonable" way to any imbalance of marketing inputs. For example, when two brands are similar in all respects (and known to be so by consumers, many of whom will have tried both), the brand with the greater weight of advertising and greater retail availability will tend to be bought more. There is no evidence that repeat-buying behaviour and the general structure of brand-choice can be influenced by factors other than perhaps real differences in product-formulation, or price, or retail availability. This is not to say that advertising, promotion and selling have no effect, but only that they influence the level of brand-shares (mostly by keeping them where they were), rather than the general structure of buyer behaviour or the "image" of the brand. In general, there is less segmentation of markets than seems often to be thought [Collins 1971].

12.4. The Way Sales Increase

The third application area of the existing theory of buyer behaviour lies in further research, as mentioned in § 12.1. It will for example be clear that the research described in this book has always built on (or used) the previous results. In the same way, further research (cf. §§ 11.6 and 11.7) should consist of practical applications of the previous findings. This point is worth stressing because research projects in marketing (and in the social sciences generally) tend often to aim at always doing something "new" or "different", instead of building on the previous results.

One example – a very important one – of the potential role of using the previous results in further research may suffice here. The problem concerns the $64,000 question of why one brand sells more than another. The research described so far has not dealt with this question. It has concerned the stationary or equilibrium situation, and whilst some of the specific applications have looked at trends, there has been as yet no *systematic* attack on studying the way in which sales increase or decrease.

Changes in people's behaviour must of necessity be more complex than their stationary or equilibrium behaviour. The question is how complex this will be – impossibly so, or just more so? Here the existence of (in effect) a single *stationary* theory makes it seem possible that the dynamic situation will itself also be amenable to attack and will produce relatively simple results. More specifically, we already know that the situation *before* any change in sales levels must be of the same general form as the situation afterwards, with only one thing differing, namely the proportion of the population buying each brand, or its market share. At the risk of some over-simplification, it follows that only three measures need to be established to characterize any particular change in sales level, namely the change in the penetration level of the brand, the time taken to effect the change, and the amount or kind of marketing effort needed to effect the change. And since the equilibrium situation of different brands are of the same general form, any such results about *changes* should also be relatively easily capable of generalisation.

Even if we look further below the surface, the situation should not (in principle) be as complex as might be thought. Thus we will want to establish the extent to which any extra buyers of the given Brand X are attracted from other brands, but the starting and end positions (of stationary brand-duplication) are already known to be simple and predictable. Again, suppose we want to try to distinguish between two kinds of changes in sales, one where a change from Brand Y to Brand X takes place in terms of certain buyers of Brand Y suddenly switching to buying Brand X, and one where there is an intermediary stage of buying *both* Brands X and Y. Then we already know that the brand-choice patterns before and after these changes are simple and of a form that is understood.

A major problem in attempting to study consumer dynamics is however that relatively few large changes in sales levels take place in real life. There are therefore relatively few situations to be studied. The classical way out is to engage in deliberate *experimentation*. Possible data-collection techniques under what may be called "semi-realistic" conditions have already been developed, such as recruiting a sample of consumers (e.g. housewives) who are offered a choice of brands each week over some extended period of time and who may be subjected to various marketing influences during that experimental period (e.g. changes in product-formulation, in price, in brand availability, in certain kinds of semi-realistic promotion, and so on). Studies involving this type of procedure have for example been reported by Tucker [1964] and McConnell [1968], and a relatively advanced technique utilizing a mobile shop for pre-test-marketing purposes has been described by Pymont [1969, 1970]. Under such conditions, it will be possible to run a large number of studies at a cost which is low relative a real-life experimentation. (See also Ehrenberg 1986, Ehrenberg and England 1987).

The analysis problem then is two-fold. Firstly, there is the sheer analysis of each experiment as such, and the integration of results from different experiments (i.e. building some integrated theory of consumer dynamics). Secondly, there is the need to translate such an experimental "wind-tunnel" theory to real-life. The application of the existing theory of stationary buyer behaviour arises here in two ways. It provides a starting-point both for analysing the experimental data as such, and, more important, for the translation process from the experimental to the real-life situation. Thus we already know what stationary buyer behaviour is like under *real-life* conditions, and the first studies of the experimental situation indicate that the same patterns tend to apply also under stationary no-trend conditions in the *experimental* type of situation [Charlton et al. 1972, Ehrenberg and Charlton 1972]. If these initial findings can be fully confirmed, then our theoretical knowledge of stationary repeat-buying and brand-choice will have facilitated a major step towards developing an experimental theory of consumer dynamics and of translating this to the real-life situation.

12.5. Summary

A variety of practical applications of the theoretical findings described in this book have already been made, but a deliberate policy of developing additional proto-type applications and systematic studies of decision-problems is needed. Buyer behaviour is a topic which is inherently relevant to marketing management, and the findings described here present a challenge: if this is what buyer behaviour is like, what implications does it have and how should we respond?

CHAPTER 13

THE DIRICHLET MODEL *

13.1. A Comprehensive Model

The Dirichlet Model brings together all the results in this book into a single and simple theoretical formulation. The model describes how frequently bought branded consumer products are purchased when the market is stationary and unsegmented, the common situation where, over the time-periods analysed,
— the sales of each brand show little variation,
— the different brands show no special groupings.
The model assumes a mixture of distributions at four levels (§ 13.2):

(i) Purchasing of the product-class takes the form of a Poisson process for each consumer,
(ii) The purchasing rates of different consumers follow 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 different consumers.

In § 13.2 we give justifications for these assumptions.

More importantly, the model has successfully described the patterns observed in more than 40 product-fields and therefore provides interpretative norms. As input the model only requires the sales level of each brand and two parameters. These can be identified as two aspects of *consumer diversity*, namely how much people differ from one another in (a) their purchasing rates and (b) their brand-choice preferences.

The model encompasses the earlier, more limited, formulations in this book which often remain easier to use. We briefly review uses of the model in § 13.4 and the more recent literature on other models in § 13.5.

*This chapter is new to the 1988 reissue. It is an edited version of a paper "The Dirichlet: A Comprehensive Model of Buying Behaviour" by Professor Gerald Goodhardt, the present author and Dr. Christopher Chatfield, read to the Royal Statistical Society in 1984 (Goodhardt et al. 1984). The mathematical notation is more general than before within this book, using suffixes j and k for two different brands rather than the x and y of Chapter 11.

13.2 The Dirichlet Model

We consider a population of N consumers making purchases in a product-class of g brands. The Dirichlet model specifies probabilistically how many purchases each consumer makes in a time-period and which brand is bought on each purchase occasion. It combines both purchase incidence and brand-choice aspects of buyer behaviour into one model. The model should therefore allow us to predict the various summary measures discussed in this book, and in Chapters 3 and 9 in particular.

In this section we give the basic assumptions and formulation of the model and empirical and theoretical justifications of the assumptions. One crucial property of the model is that different brands can simply be combined into a "super-brand". This is used technically in estimating one of the parameters of the model and in evolving theoretical predictions of specific summary measures.

The Basic Assumptions. Our formulation of the model is to specify the probability vector of the ith consumer making any specific combination $\{r_j\}$ of purchases of the $j = 1$ to g brands in an analysis-period of any chosen length T, that is r_1 purchases of brand 1, plus r_2 purchases of brand 2, etc., where $r \geqslant 0$. Summing over the $j = 1$ to g brands, $\Sigma r_j = n_i$ is the total number of purchases of the product-class made by the ith consumer in that period.

We arrive at the formulation of the particular model by making five assumptions. The first two concern brand-choice:

(A1) The ith individual's brand choices over a succession of purchases are as if random, with a probability $(p_j)_i$ of choosing brand j from $j = 1,\ldots, g$ brands. These probabilities are fixed over time and brand-choices at successive purchases are assumed independent. The number of purchases of each brand that individual i makes in a sequence of n_i purchases can therefore be modelled by a multinomial with parameters n_i, $(p_1)_i, \ldots, (p_g)_i$ [Wilks 1962, p. 139].

(A2) The probabilities $(p_j)_i$ vary among individuals according to a Dirichlet distribution (e.g. Wilks, 1962, p. 177). This is a multivariate Beta-distribution given by the joint density function

$$Cp_1^{\alpha_1-1},\ldots,p_g^{\alpha_g-1},$$

where $p_j \geqslant 0$, $\Sigma p_j = 1$, $C = \Gamma(S)/\Pi(\Gamma\alpha_j)$, $S = \Sigma\alpha_j$ and $\alpha_j > 0$. The probability of choosing brand j then has the jth marginal distribution. This is the simple Beta-distribution

$$Cp_j^{\alpha_j-1}\,(1-p_j)^{(S-\alpha_j-1)},$$

with mean α_j/S [Wilks, 1962, p. 173], which is the brand's market share. (In the traditional notation of the Beta-distribution with parameters α and β, S would equal $(\alpha + \beta)$.)

Assumptions A1 and A2 therefore say that the joint distribution of purchases of different brands across all consumers is given by a mixture of multinomials with a Dirichlet distribution. (When $g = 2$, this reduces to the well-known Beta-Binomial distribution, as noted below.)

We next make two assumptions regarding purchase incidence in the product-class:

(B1) Successive purchases of the ith individual behave as if random, and are assumed to be independent with a constant mean rate μ_i in some chosen "unit" length time-period (longer than the minimum inter-purchase time, which is usually a week for grocery products). The number of purchases n_i made in each of a succession of equal non-overlapping periods of relative length T then follows a Poisson distribution with mean $\mu_i T$.

(B2) The mean purchasing rates vary between individuals according to a Gamma-distribution with density function

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

From assumptions B1 and B2 it follows that the number of purchases of the product made by all individuals in a time-period of length T follows a Negative Binomial distribution, with mean MT and exponent K. (We here use capital letters for product-class parameters, and lower case letters for those relating to a specific brand).

Our final assumption concerns the relationship between the (A) and (B) assumptions, namely:

(C) The brand choice probabilities and the average purchase-frequencies of different consumers are distributed independently over the population.

Assumptions (A) to (C) are sufficient to specify a single model which we have called the NBD-Dirichlet model, or Dirichlet for short. In the notation of mixtures or compound distributions (e.g. Gurland, 1957), the number of purchases an individual (or household) makes of each of the g

brands in a period of length T is given by a g-variate discrete random variable with the joint frequency distribution

$$[\mathscr{M}(\mathbf{r} \mid \mathbf{p}, n) \underset{\mathbf{p}}{\wedge} \mathscr{D}(\mathbf{p} \mid \boldsymbol{\alpha})] \underset{n}{\wedge} [\mathscr{P}(n \mid \mu) \underset{\mu}{\wedge} \mathscr{G}(\mu \mid MT, K)],$$

where $\mathscr{M}$, $\mathscr{D}$, $\mathscr{P}$ and $\mathscr{G}$ denote the Multinomial, Dirichlet, Poisson and Gamma distributions.

To activate the model we need to estimate the $(g + 2)$ quantities α_j, M and K. Rather than enumerate numerically all the relevant probabilities in calculating theoretical values of any desired summary measure, we try to develop algebraic short-cuts later in this section.

Justifications of the Assumptions. 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, as documented in this book (e.g. Table 1 in the Preface) and illustrated in § 13.4. In addition there are reasons why the specific distributions of brand-choice and purchase incidence should be as just assumed and not something else.

A. The Brand-choice Distributions – Assumption A1 of a multinomial distribution is in line with the evidence outlined in Parts I and II that stochastic buying behaviour at the individual level tends to be irregular but stationary at least in the medium-term, ignoring sharp but short promotional fluctuations.

For assumption A2 there is a strong characterization result namely that given the unsegmented nature assumed for the market, the mixing distribution of the brand-choice probabilities across different individuals must be of the Dirichlet form. Thus lack of segmentation means that choosing between the different brands should in some sense be independent. But the choice probabilities for each individual are constrained to add to 1 and cannot therefore be strictly independent.

This type of problem was noted by Karl Pearson (1897). More recently, Mosimann (1962) has introduced the idea of "independence except for the constraint", whereby for two brands j and k, p_j and $p_k/(1-p_j)$ should be independent rather than p_j and p_k. This theory has been developed by Connor and Mosimann (1969), Darroch and Ratcliffe (1971), Darroch and James (1974), and James (1975), and is linked to certain ideas of rational choice by Luce (1959). It expresses mathematically what we mean in marketing when we say a market is "not segmented"; 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 now is that independence except for the constraint $\Sigma p_j = 1$ is a characterization of the Dirichlet distribution (Mosimann, 1962, 1984). The use of this distribution is therefore the quantitative analogue of our purely qualitative marketing criterion of non-segmentation. Aitchison (1982, p. 142) has said that the Dirichlet distribution seldom if ever provides an adequate description of compositional data because of its strong independence structure. But we have found that it is precisely this independence structure which empirically fits so well here. In a strictly unsegmented market where the multinomial choice probabilities for each individual are fixed over time, the Dirichlet distribution is the only possible model for brand choice.

B. The Purchase Incidence Distributions – Assumption B1 of a Poisson process with mean μ_i for the *i*th individual's purchases of the product-class over time, 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) and so irregular that it can be regarded *as if random.* The Poisson assumption and possible alternatives, like some Erlang distribution for interpurchase times, have been extensively discussed, especially for brand purchases (e.g. Ehrenberg, 1959a; Herniter, 1971; Chatfield and Goodhardt, 1973; Dunn et al., 1983; and in this book). The Poisson process remains a workable approximation.

Assumption B2 of a Gamma mixing distribution for the values of μ_i can probably be justified as follows. If for different product-classes *P, Q, R, S,* etc (like toothpaste, breakfast cereals, canned soup, etc.)

(1) the average purchase rate of *P* is independent of the rates for the other products *Q, R, S,...,* and

(2) *P*'s proportion of a consumer's total purchases, namely $P/(P + Q + R + S + \ldots)$, is independent of her total rate of purchasing all the products,

then it can be shown, following a similar characterization for brands (e.g. Goodhardt and Chatfield, 1973; Chatfield, 1975), that the distribution of the mean rates of purchase of *P* must be Gamma. These independence conditions are likely to be approximately fulfilled in practice. Thus heavy buyers of toothpaste are not necessarily heavy buyers of canned soup, nor do heavy buyers of the whole range of products necessarily devote an above average proportion of their consumer goods expenditure to toothpaste. In both cases, any correlation – e.g. due to household size – is likely to be fairly low. No direct empirical analyses for different product-classes have yet been made, but much suitable data exists.

Assumption C is in line with general experience across a wide range of product-fields, including the specific cases referred to in Table 13.4. Shoemaker et al. (1977) have noted that in three product-fields some of the brands had varying shares among light, medium and heavy buyers, but the differences appear to have been small and not consistent.

The Additivity Property. An important property of the model is that any two brands j and k with means a_j/S and a_k/S can be combined into a super-brand with mean $(a_j + a_j)/S$. Nothing else in the specification of the Dirichlet model is affected [Wilks, 1962]. This feature is not common to other such models (e.g. Herniter, 1971).

The additivity property helps to explain why minor brands can be successfully grouped into an "all other brands" category and, more generally, why the model can cope when a brand is made up of different pack-sizes or flavours, or is bought from different retail outlets. It is also of marked computational help in estimating the model and dominates the algebraic short-cuts we use in calculating values of summary measures.

13.3 Fitting the Model

In order to apply the Dirichlet model to a g-brand market, we must first estimate the values $\alpha_1, \alpha_2, \ldots, \alpha_g$ of Assumption A2 and the parameters M and K of Assumption B2. In principle we assume that detailed data are available for some "unit length" base period (say 12 weeks). In practice, we use summary statistics, namely penetrations and average purchase rates, to fit the model. In fitting the Dirichlet to the full observed data see Nelson (1986).

Using the g observed per capita (or per household) average purchasing rates m_j as input variable, we can calculate the product-class purchasing rate M as $M = \Sigma m_j$, and equate the theoretical and observed market-shares (the means of the marginal Beta-distributions)

$$\alpha_j/\Sigma\alpha_j = m_j/\Sigma m_j.$$

As the brand shares must add up to one, there are only $(g-1)$ independent equations of this form. To solve the equations we need to estimate $\Sigma\alpha_j$, which we call S – one of the two structural parameters of the model.

We believe that a simple and effective way to estimate S is to equate the observed proportion $1 - b_j$ of the population not buying brand j in the chosen time-period to the corresponding theoretical probability, and solve for S, and then form a weighted average of the separate values of S across all g brands. The theoretical Dirichlet formula for $(1 - b_j)$ is however not in closed form for S and hence needs an iterative solution. It also involves an infinite series for the NBD probabilities of buying the product-class and this calls for an ad hoc truncation rule. The details are set out in Appendix C.

The parameter K is calculated by fitting an NBD to the distribution of purchases for the whole product-class. If the distribution is reverse J-shaped, we fit by "mean and zeros", i.e. by solving $1 - B = (1 + M/K)^{-K}$ for K. This then has high efficiency (e.g. Anscombe, 1950; Ehrenberg, 1959a; Chatfield, 1969). If the distribution is not reverse J-shaped, we use the method of moments. The theoretical NBD of product-class purchases can then 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 = 0, 1, 2,... purchases and compared with the observed distribution in order to assess goodness-of-fit.

The Empirical-Dirichlet Model. For some products the distribution of product-class purchases shows systematic deviations from an NBD. Toothpaste, which we use illustratively in § 13.5, is an example. Table 13.1 shows that there is a small short-fall of once-only buyers (19% versus 22%). We believe that such discrepancies occur for "saturated" markets, i.e. products which generally have no direct substitute like toothpaste, and where the additivity type of property then does not apply at the product-class level. But more extensive empirical work is needed.

Table 13.1 Purchases of Toothpaste in 12 Weeks

(Fitted by mean and zeros)

	0	1	2	3	4	5	6+	*Av. no. per buyer*
Observed (%)	44†	19	14	9	6	3	4	2.6†
NBD (%)		22	13	8	5	3	5	

†Used in fitting.

The departure from an NBD here is not large enough to make us discard the NBD part of the model. But where it is, we can still use the rest of the

Dirichlet model by taking the observed distribution of product-class purchases as input instead of the fitted NBD. Thus we replace the theoretical NBD probability P_n in the various formulae by the observed proportion of consumers making n product purchases. The model is then referred to as the Empirical-Dirichlet model. It tends to give a better fit in the base period, but is less elegant and we are no longer able to make predictions for time-periods of a different length T. Thus we normally prefer to use the full NBD-Dirichlet model except when the goodness-of-fit is rather poorer than in Table 13.1.

Applying the Model. Having estimated the parameters, we can calculate the theoretical value of any specific aspect of buying behaviour, such as the summary measures discussed earlier in this book. We can do this in the base-period and also for any period of length T ($T \gtreqless 1$). The calculations are straightforward in principle though computationally somewhat tedious, even when using algebraic short-cuts.

Our main simplification is to try to reduce the calculations down to those of a Beta-Binomial distribution (e.g. Chatfield and Goodhardt, 1970). Thus in calculations for a specific brand j with brand-share a_j/S, we combine all the other brands into to a single superbrand with brand-share $(S - a_j)/S$. This means that the probability of making r_j purchases of brand j, conditional on n purchases of the product-class having been made ($r_j \leqslant n$), is given by the Beta-Binomial distribution

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

where B here denotes the Beta function.

The proportion of consumers buying the product-class n times and buying brand j r_j times is then given by the product of the equations for P_n and $p(r_j \mid n)$

$$p\,(r_j,\ n) = P_n\, p(r_j \mid n).$$

(A capital P stands for probabilities of buying the product-class, a lower case p for those of buying a brand.) By summing $p(r_j,\ n)$ over appropriate values of n and r_j, we can calculate any statistic of interest for the brand. Summations over the values of n need to be truncated, for example in the way described in the Appendix C for estimating S. In these computations both P_n and $p(r \mid n)$ can be calculated using appropriate recurrence relations.

Estimates for Brand j. Examples for brand j (dropping the subscript j) are: The penetration b of brand j is estimated as $1 - p(0)$, the proportion not buying the brand, where

$$p(0) = \sum_{n=0} \{P_n \, p(0 \mid n)\}$$

and $p(0 \mid n)$ is the probability of making zero purchases of brand j given that n purchases of the product-class have been made in the analysis-period. The summation is again truncated.

Here $p\,(0 \mid n)$ is, from the Beta-Binomial, writing α for α_j,

$$p(0 \mid n) = \frac{(S-\alpha)\,(S-\alpha+1)\ldots(S-\alpha+n-1)}{S(S+1)\ldots(S+n-1)} \quad \text{for } n \geqslant 1$$

$$p(0 \mid 0) = 1$$

The theoretical number of purchases of brand j per buyer is calculated as

$$w = \sum_{n=1} \left\{ P_n \sum_{r=1}^{n} rp(r|n) \right\} \Big/ [1 - p\,(0)],$$

and their average number of purchases of the product-class as

$$w_P = \sum_{n=1} \{nP_n\,[1 - p(0|n)]\}/[1 - p(0)].$$

The proportion of the population who buy brand j only (the "sole" buyers) is given by a computationally very effective short-cut

$$\sum_{n=1} \{P_n p(n|n)\},$$

since if they buy the product n times they must be buying the brand n times also. Their average purchase frequency per buyer is

$$\left\{ \sum_{n=1} [nP_n p(n|n)] \right\} \Big/ \left\{ \sum_{n=1} [P_n p(n|n)] \right\}.$$

In a period of length T, relative to the base period (with $T \gtrsim 1$), all the above formulae are unchanged in the NBD-Dirichlet except that M in the NBD equation becomes MT.

To estimate the number of repeat buyers from one period of length T to another of equal length we enumerate b_{2T} for the double period and b_T for a single period and calculate $2b_T - b_{2T}$. But we have at this stage developed no simple way of calculating the theoretical average purchase frequency of such repeat-buyers in each period. Instead, we generally use the more convenient NBD/LSD theory in this area.

Duplicated buyers of brands j and k. Theoretical estimates of purchase combinations of two or more specific brands, *j, k,* etc, would generally require more number-crunching than we have yet attempted. (Another approach would be to use the model to construct theoretical purchases of a "sample" panel of 10,000 simulated households and then tabulate the relevant summary measures, as we routinely do with observed data.)

A simple method has however developed for estimating the number of duplicated buyers of any particular pair of brands j and k. We use the additivity property of § 13.2 in a different way. We form a composite brand $(j + k)$ and estimate its penetration $b_{(j+k)}$. This then gives us b_{jk}, the theoretical proportion of the population buying both brands at least once as

$$b_{jk} = b_j + b_k - b_{(j+k)},$$

and hence also "conditional" proportions like $b_{j/k} = b_{jk}/b_k$.

13.4 The Fit of the Model

The Dirichlet model directly or indirectly describes the buying patterns that have been found for many different products, food and non-food, over many years, in the UK and USA, etc. Thus in the relevant parameter range, the Dirichlet gives much the same results as the successful earlier sub-models in this book.

Here we illustrate the fit of the model with results for toothpaste in the UK [Aske Research, 1975]. The input is in effect that in the January to March quarter of 1973, 56% of the AGB panel of 5240 continuously-reporting households bought toothpaste on average 2.6 times (Table 13.1) which gives $M = 1.5$ and $K = .78$, and that the brand penetrations averaged 9% to give $S = 1.2$. Table 13.2 shows how, given the market-shares of the eight leading brands, the model recovers the individual brand penetrations b, and purchase frequencies w and w_P for the quarter and, extrapolating, for the year as a whole (i.e. putting $T = 4$).

Table 13.2. Penetration and Average Purchase Frequencies of Toothpaste Brands

(Observed "O" and Dirichlet "D", Fitted to Quarterly Data)

		Quarterly						Annual					
		%		Av. Purchases of				%		Av. purchases of			
Leading brands (and market-shares)		*Buying*		*Brand*		*Product*		*Buying*		*Brand*		*Product*	
		O	D	O	D	O	D	O	D	O	D	O	D
Colgate DC	(25%)†	20	20	1.8	1.8	3.1	3.2	34	37	3.7	3.8	10	8.7
Macleans	(19%)†	17	17	1.7	1.7	3.0	3.3	32	32	3.2	3.6	10	8.9
Close Up	(10%)†	9	8	1.6	1.7	3.5	3.3	15	17	3.0	3.2	10	9.2
Signal	(10%)†	8	9	1.9	1.7	3.5	3.3	17	18	3.4	3.3	11	9.2
Ultrabrite	(9%)†	8	8	1.7	1.7	3.2	3.3	17	17	2.9	3.2	10	9.2
Gibbs SR	(8%)†	7	7	1.7	1.7	3.2	3.3	17	14	2.8	3.2	10	9.3
Boots Priv. Label	(3%)†	3	2	1.4	1.7	2.6	3.4	6	5	2.4	3.0	9	9.4
Sainsbury Priv. Label	(2%)†	2	2	1.5	1.6	3.1	3.4	3	4	3.2	3.0	11	9.5
Average		9†	9	1.7	1.7	3.2	3.3	18	18	3.1	3.3	10	9.2

[†Used in fitting]

For example, for Beecham's Macleans with a market-share of 19% the model predicts:

(a) That in the quarter, $b = 17\%$ of households bought it on average $w = 1.7$ times and made a total of $w_P = 3.3$ purchases of any toothpaste.

(b) That in the year as a whole, $b = 32\%$, $w = 3.6$ and $w_P = 8.9$.

The fit is close, e.g. to within about ± 1 for the bs across all eight brands. The only bias concerns w_P in the year, with $\overline{\text{O}} = 10$ and $\overline{\text{D}} = 9$. This seems due to the short-fall of once-only buyers in Table 13.1. Fitting the Empirical-Dirichlet to the annual data gives $\overline{w}_P = 10$, as observed.

The predictions also reflect the usual small Double Jeopardy and Natural Monopoly trends in w and w_P noted in Chapters 3 and 9, namely that

(c) The w's generally decrease with market-share. (The observed values jump around somewhat with subsample sizes down to 100. But the four largest brands average $\bar{O} = 3.3$ and $\bar{D} = 3.5$ in the year, and the four smallest $\bar{O} = 2.8$ and $\bar{D} = 3.1$.)

(d) There is a contrary small *upward* trend in the theoretical values of w_P though not in the observed ones here. It has however been observed in almost all the other cases analysed (e.g. Chapters 9 and 10 here; Ehrenberg and Goodhardt, 1976; Ehrenberg and Goodhardt, 1979a; Wrigley, 1980; Kau, 1981, and numerous Aske Research reports). It occurs as a statistical selection effect despite Assumption C and the independence structure of the model.

The Dirichlet model also successfully predicts the other measures of buying behaviour, such as

(e) The frequency distributions of purchases of the individual brands, as illustrated in Table 13.3.

Table 13.3. Distribution of Annual Purchases of Toothpaste

(Two typical brands)

		Number of purchases						
48 weeks		0	1	2	3	4	5	6+
Macleans	Observed	67	13	6	4	2	2	5
	Theoretical	68	12	6	4	3	2	5
Sainsburys	Observed	94	3	1	1	0.5	0.3	0.4
	Theoretical	94	3	1	1	0.4	0.2	0.6
Average†	Observed	81	8	4	2	1	1	3
	Theoretical	81	8	4	2	1	1	3

†The eight leading brands.

(f) The penetration and average purchase frequency among frequent or infrequent buyers of the product-class, as shown in part in Table 13.4 (a rather direct test of the independence Assumption C).

(g) The incidence of 100% loyal or "sole" buyers of each brand, to within a few per cent both in short and long period (not shown here in detail). There is a suggestion in the limited data analysed so far that the Dirichlet somewhat underestimates the average purchase frequencies of sole buyers. More work is needed here.

(h) The proportion of the quarterly buyers of Macleans who also buy a specific other brand at least once in the quarter, as shown in Table

Table 13.4. Penetration and Purchase Frequency among Infrequent (1–6) and Very Frequent (13+) Buyers of Toothpaste

	Annual purchases of toothpaste							
	1–6 purchases				13 + purchases			
	b		*w*		*b*		*w*	
48 weeks	*Obs.* %	*Theo.* %	*Obs.*	*Theo.*	*Obs.* %	*Theo.* %	*Obs.*	*Theo.*
Colgate DC	35	37	1.8	1.8	65	68	6.3	6.2
Macleans	33	32	1.7	1.8	64	61	5.4	5.8
Close Up	24	24	1.7	1.7	54	44	4.7	5.2
Signal	14	14	1.8	1.7	44	40	5.3	4.9
Ultrabrite	17	16	1.6	1.6	35	35	4.7	4.7
Gibbs SR	17	16	1.8	1.6	34	34	4.4	4.7
Boots Priv. Label	6	5	1.9	1.6	11	12	3.0†	4.1
Sainsbury Priv. Label	2	3	1.9	1.5	8	8	4.5	4.1
Average	19	19	1.8	1.7	39	38	4.8	5.0

[† Sub-sample base: 81 households]

13.5. The annual predictions here show a 10–20% bias in line with that for w_P in Table 13.2 and also reflect some discrepancies in the trend with T of $b_{j/k}/b_j$ (Aske Research, 1975a gives details.) A related problem is that the parameter S in the model should be constant irrespective of T, but it varied from 1.0 when estimated from 4-week data, to 2.2 for the 12-week data, to 1.8 in the year. This needs further study. Nonetheless, the predictions illustrated in (a) to (h) generally work well both here and in many other product-fields, as illustrated further in Table 13.7 in the next section.

Table 13.5 Quarterly Duplication of Macleans and the Other Brands

(% Buyers of Macleans also buying Other Brand)

	Colg. DC	*Macl.*	*Close Up*	*Signal*	*Ultra-brite*	*Gibbs SR*	*Boots PL*	*Sains. PL*
Observed (%)	28	(100)	10	9	9	10	2	1
Theoretical (%)	24	(100)	11	10	10	8	4	2

13.5 Uses of the Model

The Dirichlet model summarizes many aspects of buying behaviour. Practical applications include providing interpretative norms, both for stationary and non-stationary markets. This is in line with the discussions of Chapters 5 and 6, but a few further illustrations are probably worthwhile.

Providing norms for stationary markets. Suppose some data show that 53% of the buyers of Macleans toothpaste in a given quarter bought it again in the next quarter. The question is whether this is low (*only* 53% repeat-buyers), or high (*as many as* 53%), or what? The model's prediction of 51% then tells us that the observed rate is neither particularly low nor high, but merely just about normal for any toothpaste brand with a 19% market-share. (Similarly, had the observed figure been 65%, we could judge it *high* only by having an interpretative norm.)

Such interpretative norms are needed whenever we deal with a new data set. For example, some years ago we were faced for the first time with a range of data from the USA. But we did not have to start from scratch: the earlier NBD model of repeat-buying had already been well established in the UK and was then found to apply to the US data as well. This could be summarized succinctly (Ehrenberg and Goodhardt, 1968b) as "American and British repeat-buying habits are the same", a finding which has since been extensively confirmed (e.g. Aske Research, 1969, 1974, 1975, Uncles and Ehrenberg 1987b).

Recent extensions to consumers' *store-choice* (e.g. Wrigley, 1980; Kau and Ehrenberg, 1984) could similarly be summarized as "Store-choice is like brand-choice". Table 13.6 illustrates in detail how the penetration growth of instant coffee at different store-groups can be closely predicted by the same Dirichlet or NBD models as used for brands (Kau, 1981).

Table 13.7 similarly shows the general fit of the model for half-yearly duplication of purchasing between the different store-groups. The model predicts that the duplication in each column should be about the same, effectively proportional to the brand penetrations (as in the old Duplication of Purchase Law – see Chapter 10), but with a very small upward trend with decreasing market-share, as already noted empirically. The fit for the individual duplications in Table 13.7 is to within a mean deviation of 4 percentage points, representing a .91 correlation between the observed and predicted values.

A further illustration of the use of interpretative norms is that 10 years ago it was found that almost two-thirds of the customers of the London

Table 13.6 Penetration Growth of Instant Coffee at Different Store-groups
(Four leading store-groups)

		Analysis-period (in weeks)			
		1 week	*6 weeks*	*12 weeks†*	*24 weeks*
Any instant coffee	Obs. (%):	17.0	54	65	76
	Theo. (%):	16.0	50	65†	77
Coop	Obs. (%):	3.8	16	20	28
	Theo. (%):	3.7	14	20†	27
Tesco	Obs. (%):	2.2	9	14	21
	Theo. (%):	2.0	9	14†	20
Independents	Obs. (%):	0.9	5	7	11
	Theo. (%):	1.0	4	7†	10
Fine Fare	Obs. (%):	0.4	3	3	5
	Theo. (%):	0.5	2	3†	5
Av. store group‡	Obs. (%):	2.1	9	13	18
	Theo. (%):	2.1	8	13†	17

[‡ Eight leading groups.] [† Used in fitting.]

Table 13.7. Duplication of Purchase of Different Store-groups
(Instant Coffee)

		% also buying at							
24 weeks		*Misc.*	*Coop*	*Kwik*	*Tesco*	*Asda*	*Indep*	*Symb*	*Fine*
Buyers at									
Miscellaneous	%	–	32	25	29	13	15	16	8
Coop	%	38	–	21	22	19	13	12	4
Kwiksave	%	37	26	–	22	9	11	9	7
Tesco	%	45	29	23	–	17	13	14	7
Asda	%	31	38	14	24	–	9	14	3
Independent	%	52	34	23	24	12	–	18	5
Symbol	%	52	32	20	30	20	19	–	8
Fine Fare	%	50	23	29	27	9	11	16	–
Average	%	43	31	22	25	14	13	14	6
1.2 × Penetration†	%	40	34	26	25	17	13	12	6
Penetration	%	33	28	22	21	14	11	10	5

†Approximate Dirichlet prediction.

Business School's executive courses had sent only a single participant in 2 years. Should this have been a cause for concern? The close fit of the NBD or Dirichlet [Charlton and Ehrenberg, 1976a; Powell and Westwood, 1978; Sichel, 1982] showed that the pattern was little or no different from that for brands of breakfast cereals, detergents, toothpaste, petrol, or whatever. Similar findings arise with airline contracts for aviation fuel with oil companies across different airports (Ehrenberg, 1975), and with a consulting firm's clients (Coggill and Simpson, 1984).

More general conclusions include that advertising is not needed to create the normal levels of brand loyalty (Charlton and Ehrenberg 1976b) and that buyer behaviour can be normal immediately after a major upset, with no lagged effects (Ehrenberg and Goodhardt, 1969; Ehrenberg 1984b).

Interpreting change. Theoretical norms for a stationary market also provide a base-line for interpreting *change* (i.e. non-stationary situations), without having to match the results against an empirical "control sample". Thus just by comparing non-stationary data with the stationary norms we can assess whether an increase in sales came from attracting more buyers or from existing buyers buying more, and if the latter, whether it was heavy or light buyers doing so. Analyses of real life cases have been discussed elsewhere (e.g. in Part III of this book, and more recently in Ehrenberg and Goodhardt, 1979a, and Wellan, 1987a & 1987b) as well as examples involving deliberate experimentation (e.g. Charlton and Ehrenberg, 1976a & 1976b; Ehrenberg, 1981, 1986; Ehrenberg and England, 1987; Motes and Woodside, 1984; Motes et al., 1984; Castleberry, 1983, and Castleberry et al., 1987).

Prescriptive uses. The model can also help in decision-making. Table 13.2 implies that a brand like Ultrabrite could double its sales by getting more buyers, since other brands have higher *b*s. It might also seem that they could buy more Ultrabrite without having to use more toothpaste, since Ultrabrite's customers mostly buy other brands as well ($w_P = 10$ and $w = 3$ in the year). But the model tells us that this will not happen and it is supported by the facts: w hardly differs for the different brands, whether observed or predicted. A marketing plan which aims at increasing sales by getting existing buyers to buy much more of the brand would therefore be aiming at something altogether unusual or unlikely, like making pigs fly. There is a law of nature against it, or so it would appear. A substantial application in NPD (New Product Development) is given elsewhere (Ehrenberg, 1987).

Understanding the nature of markets. The generalized empirical knowledge summarized by the model also helps us to understand better the nature of consumer goods markets. Examples are

- The number of customers which a brand (or retail store) has varies dramatically, but predictably, with the length of the time-period analysed.
- Brand loyalty exists but is low and not exclusive (i.e. consumers habitually or "loyally" buy more than one brand).
- Different aspects of loyalty, like repeat-buying and multi-brand buying, are directly related.
- Brands hardly differ from each other in the degree of loyalty each attracts. In the medium term, competitive marketing inputs (pricing, advertising, etc.) only show up in the brands' market-shares.
- Most branded goods markets are largely unsegmented.
- The structure of buyer behaviour is the same for radically different kinds of productclasses (like breakfast cereals or detergents), for different advertised brands (like Kellogg's Corn Flakes and Nabisco's Shredded Wheat), and irrespective of "exogenous" variables like the interest rate, or that the toothpaste data in § 3 arose 2 years after the 1973 oil crisis, and while Mr Ford was President of the USA.

The Nature of the Input. The input required by the Dirichlet model is not only very parsimonious — the g values of a_j and the two structural parameters S and K — but is also readily interpretable.

The sales levels. As already noted, the g values of a_j reflect the per capita or per household purchasing levels m_j of the g brands in the chosen unit time-period. A possible measurement simplification is that these values could be obtained from sales data (e.g. retail audits or ex-factory shipments), without having to observe individual consumers' purchases at all.

In practice we reformulate the m_j as $M = \Sigma m_j$, the average rate of purchase of the total product-class and the market shares m_j/M (which amount to $(g-1)$ independent values). With this formulation only M varies with T, the length of the time-period analysed.

The different brands in a product-class generally differ and compete with each other in many ways (product-formulation, pricing, packaging, advertising and promotion, retail distribution, usage patterns perhaps, etc.). But the close fit of the model shows that in an unsegmented near-stationary market all these other variables have generally no net effect on the structure of buyer behaviour, but are subsumed by the brands' market-shares. Such a strong causal interpretation is possible because it is negative (see also Ehrenberg 1982, Chapter 20). We are saying that lack of correlation probably implies lack of causation.

Two measures of diversity. The two parameters, K and S, are characteristics of the product-class and can be interpreted as reflecting the heterogeneity or diversity of consumers. Thus from Assumptions B1 and B2 in § 13.2, the standard deviation of the resultant NBD is given by $\sqrt{\{M(1 + M/K)\}}$. Hence K reflects how much peoples' individual product-purchases differ from the overall mean M.

In Assumptions A1 and A2, the variance of the individual probabilities $(p_j)_i$ in the marginal Beta-distribution for brand j is $\{m_j\ (M - m_j)\}/\{M^2(1 + S)\}$. S therefore reflects the extent that people differ from each other in their propensities to buy each brand. Thus if S is very large, the variance is near-zero and everyone has much the same probability m_j/M of buying brand j (minimum diversity). If S is small, the individual $(p_j)_i$ differ more across i. Indeed, the distribution reduces to just two "spikes" if $S = 0$: a proportion m_j/M of individuals that always buy brand j and a proportion $(1 - m_j/M)$ that never do. (This is maximum diversity; but no one switches between brands.)

Neither of these two kinds of consumer variability need however be measured directly. Instead we can in principle, and often in practice, estimate K from the observed values of M (which could be got from sales data, as already noted) and $(1 - B)$, the observed proportion of non-buyers of the product. The latter is a qualitative rather than a quantitative form of input. We can similarly estimate S without using any quantitative measurement of peoples' brand-choice behaviour. In effect we use only $(1 - \bar{b})$, a weighted average of the proportions of the population who do *not* buy each brand in principle. One therefore does not need the individual details of continuous consumer panel data to activate the model.

13.6 The Background

In this section we briefly refer to our more limited models of buyer behaviour earlier in the book and also give a brief up-date on some other models of buying behaviour.

Our Previous Models. The Dirichlet model describes stochastically when a purchase of the product-class is made and which brand is chosen. In contrast, our previous modelling started with when an individual brand was bought, using the Poisson-Gamma NBD model as in Assumptions B1 and B2, but applied to the individual brand, not the product-class. Finding various generalizable empirical regularities and developing theoretical formulations then generally went hand in hand. Next, we linked the results for different brands, for example that b and w for different brands j and k could be related as $w_j(1-b_j) \doteqdot w_k(1-b_k)$ = constant. Finally, there were results concerning the extent to which any one consumer bought more than one brand over time. Included here were results showing that the proportion of buyers of brand j who only buy j in the analysis-period should, on an independence assumption, equal $(1 - B)/(1 - b_j)$, and that brand duplication between brands j and k follows the Duplication of Purchase Law $b_{j/k} = Db_j$, where D is constant for all pairs of brands.

The derivation and generally close fit of such sub-models of specific aspects of buying behaviour has been described earlier in some detail. The Duplication of Purchase Law and the level of w_P were, however, two empirically-established regularities for which no predictive theoretical models had been developed, but which are now also accounted for by the Dirichlet. Further extensions and applications of some of these earlier models were also developed by others (e.g. Rothman (S.R.S, 1965); Grahn, 1969; Morrison, 1969; Charlton et al., 1972; Jephcott, 1972; Charlton and Ehrenberg, 1976a; Paull, 1978; Easton, 1979; Frisbie, 1980; Wrigley, 1980; Greene, 1982; Sichel, 1982; Wrigley and Dunn 1984a,b).

The characterization of the Gamma-distribution in the NBD model for brands [Goodhardt and Chatfield, 1973] then led to the development of a special case of the Dirichlet model with complete independence between different brands as a direct generalization of the NBD model for specific brands [Aske Research, 1973; Chatfield and Goodhardt, 1975; Chatfield, 1975]. This in turn led to the more general but mathematically different Dirichlet model (e.g. Aske Research, 1974, 1975; Ehrenberg and Goodhardt, 1976).

Our earlier models give numerically predictions close to those of the Dirichlet model, in the usual parameter range of the observed data. However, they are mathematically different. For example, the distribution of purchases of a given brand in the Dirichlet model is not an NBD (except in the case of independence, when $S = K$). Nonetheless, the two tend to agree closely as illustrated in Table 13.8; in general they agree if anything more closely with each other than either does with the observed data. The older sub-models like the NBD still tend to be used in practice because predictions for a single brand are computationally simpler and are not affected by non-stationarities for the other brands.

Table 13.8 The Frequency Distribution of Brand Purchases

(Observed, Empirical-Dirichlet and NBD)

		Number of purchases						
48 weeks		0	1	2	3	4	5	6+
Macleans	Obs. (%)	67	13	6	4	2	2	5
	Dir. (%)	68	12	6	4	3	2	5
	NBD (%)	(67)†	13	6	4	3	2	5
Average brand	Obs. (%)	81	8	4	2	1	1	3
	Dir. (%)	81	8	4	2	1	1	3
	NBD (%)	(81)†	8	4	2	1	1	3

†Used in fitting.

Other Models. Other published work on mathematical models of buyer behaviour has mainly been concerned with how markets change rather than with first describing the steady-state case (see Massy et al., 1970; Montgomery and Ryan, 1974; Lilien and Kotler, 1983, for reviews). The first-order Markov model (see also p. 214) remains the best known example. As stated in Chapter 11, it assumes that the probabilities of repeat-buying and of switching from brand j to k for successive purchases are invariant characteristics of each brand, i.e. remain constant even when the market-shares of the brands change. But empirically the reverse is the case. Repeat-buying and switching propensities tend to be independent of the specific brands and vary instead with their market-shares. The known facts and the Markov concept could not differ more (see also Ehrenberg, 1965c) and Markov seems slowly to be dropping out of favour, although not yet fully.

A formulation which may be dealing with stationary markets is the Hendry model (Butler, 1966; Herniter, 1973; Kalwani and Morrison, 1977; Ehrenberg and Goodhardt, 1979b). It centres on the equivalent of the Duplication of Purchase Law $b_{j/k} = Db_j$, but only for pairs of successive purchases (where D becomes $S/(S + 1)$ in terms of the Dirichlet parameter). This ignores purchase incidence and population heterogeneity (as does the Markov model). In addition, the model assumes that D can be estimated from a rather incoherent "entropy" assumption (e.g. Ehrenberg and Goodhardt, 1973), which does not in itself lead to a good empirical fit (e.g. Ehrenberg and Goodhardt, 1974).

A much more substantial approach to stationary markets is the work by Bass and his colleagues (e.g. Bass et al., 1976; Jeuland et al., 1980) who formulated models similar to the ones here, almost certainly starting in part from Chatfield and Goodhardt (1975). But the work has lacked systematic evidence of empirical regularities to be fitted by the model. It also did not give characterizations of the distributions as in § 13.2, nor an adequate estimation procedure for the crucial parameter S. More generally, the Dirichlet has been gaining increasing currency in recent discussions of consumer behaviour models as a convenient extension of the univariate beta-distribution (e.g. Rust, 1986). An early mention was Pyatt (1969). Among discussions linking such stochastic models to so-called explanatory variables the work of Zufryden is particularly noteworthy (e.g. Jones and Zufryden, 1980. See also Uncles, 1987).

Further Work. Although the Dirichlet model now provides us with a good tool for handling stationary purchasing data, a good deal of further work still needs to be done. The earlier tables showed some discrepancies from the model. They are mostly small and may be of little marketing importance. But are they generalizable and, if so, can they then be related to other variables, like pricing or advertising or future market trends?

Other discrepancies concern the lack of constancy of the parameter S and similar problems for the trend of the duplication coefficient D with T, and the cases where the NBD does not fully fit (the "Empirical-Dirichlet" situation of Table 13.1). There is also at least one product-category, cigarettes, for which repeat-buying appears to differ somewhat from the NBD/Dirichlet predictions (Aske Research 1971).

Analytically, developing closed-form approximations to the Dirichlet would be very helpful especially where NBD/LSD approximations do not already exist (e.g. for the duplication coefficient D or for w_P). We also need to explore the numerical nature of the model more fully.

An empirical question is whether the two basic "diversity" parameters K and S are product-class characteristics which are the same for different demographic sub-groups, different store-groups, different countries, and so on.

Extensions to other market conditions, types of economies and types of products beyond frequently bought branded consumer goods are also needed. Some have already been started, as noted in § 13.3. But more can be done, for example with phenomena where the NBD, the Duplication Law or at least related types of patterns are already known to occur. This includes people's viewing of television programmes (e.g. Goodhardt et al., 1975), the readership of print media (e.g. Hyett, 1958; Leckenby and Kishi, 1984), the incidence of human accidents (e.g. Irwin, 1964) and the original NBD studies on the abundance of species (e.g. Fisher et al., 1943).

We can also increase the number of choice dimensions from one – i.e. which brand of a product is chosen, or at which store the product is bought – to cover both brand and store choice [Kau and Ehrenberg, 1984], or brand and flavour, brand and pack-size, etc. The choices might possibly be hierarchical, e.g. first brand then flavour, or vice versa (Goodhardt in discussion on Aitchison, 1982).

An altogether major extension will be to model non-stationary markets. But here we first need many empirical studies to see what if any generalizable patterns exist (e.g. Ehrenberg 1981; and Wellan, 1987a,b; Ehrenberg and England 1987).

More generally still, there is also scope for using the model to help our understanding of competition, of marketing inputs like advertising, price, product-quality and retail distribution, and of the psychologically "low-involvement" nature of consumers' brand-choice (e.g. Barwise, 1984).