HARRIS SCHOOL WORKING PAPER
SERIES 04.1

Market Shares: Some Power Law Results and Observations

Rajeev Kohli and Raaj Sah

MARKET SHARES: SOME POWER LAW RESULTS AND OBSERVATIONS

Rajeev Kohli
Columbia University

Raaj Sah*
University of Chicago

Revised: November 21, 2003

*We thank Petr Barton, Sunil Gupta, Kamel Jedidi, Don Lehmann, Puneet Manchanda, Irina Ostrovnaya, and Suchandra Thapa for their help. The paper has benefited from numerous comments and suggestions from seminar participants at Columbia University, Harvard University and the University of Pittsburgh.

MARKET SHARES: SOME POWER LAW RESULTS AND OBSERVATIONS

Abstract

The authors report an empirical regularity in the market shares of brands. They propose a theoretical explanation for this observed regularity and identify additional consequences of their explanation. Empirical testing of these consequences supports the proposed explanation.

The empirical regularity is obtained using cross-sectional data on market shares of brands in 91 product categories of foods and sporting goods sold in the United States. In total, the data set has 1171 brands. The key feature of the empirical regularity is that, in each category, the difference between pairs of successively-ranked market shares forms a decreasing series. In other words, the decrease in the market share between two successively-ranked brands becomes smaller as one progresses from higher-ranked to lower-ranked brands. The observed patterns of market shares are represented by the power law to a remarkable degree of accuracy, far surpassing the fits typically observed in predictive models of market shares.

The authors then propose an explanation for the regularity in terms of a model of consumer purchases. At the heart of this model is a special case of the well-known Dirichlet-multinomial model of brand purchases. This special case corresponds to Bose-Einstein statistics, and has an intuitive implication that the purchase probability of a brand is proportional to the number of units of that brand already purchased. A notable feature of the model is that it considers multiple product categories, and allows for the entry and exit of brands over time. The authors then describe, and test using the data set, two predictions of the proposed model regarding patterns of market shares across product categories. For both of these predictions, which appear counterintuitive, the authors report reasonable empirical support. Overall, there appears to be good evidence supporting the proposed model and the observed regularities. The authors discuss some potential implications of the regularities for marketing practice and research. They also offer an interpretation of the previously known square-root relationship between market share and the order of entry of firms into an industry.

I. INTRODUCTION

Bass (1995, p. G7) defines an empirical generalization as "a pattern or regularity that repeats over different circumstances and that can be described simply by mathematical, graphic, or symbolic methods." Unlike a law of classical physics, an empirical generalization does not have to be universal over all circumstances, and it is not necessary for the parameters governing the regularity to be invariant across different circumstances. Such empirical generalizations are typically approximate rather than exact, and they are descriptive rather than directly causal (Ehrenberg 1982). These generalizations facilitate the important task of constructing theories (Ehrenberg 1995) and of generating testable consequences beyond the original data (Simon 1968). The Pareto income distribution is an enduring empirical generalization in the social sciences (Persky 1992). In marketing, the Bass diffusion model (Bass 1969) and the Dirichlet model of brand purchases (Chatfield and Goodhardt 1975, Goodhardt, Ehrenberg and Chatfield 1984) are important empirical generalizations.

We present a number of empirical findings concerning patterns of market shares of brands. We offer a consumer-purchase model for explaining the observed patterns, and present qualitative properties of this model that appear especially appealing for the study of market shares. We test two seemingly counterintuitive predictions of this model, and find that there is good empirical support for them. We also present some implications, as well as several qualitative observations, that arise from or are related to the findings reported in the paper.

Most of the empirical findings presented in this paper are, to our knowledge, unavailable in the literature. A notable feature of the proposed explanatory model is that it is connected to the Dirichlet-multinomial model of brand purchases. This is significant because, as Uncles, Ehrenberg and Hammond (1995) note, the Dirichlet model may be the best-known empirical generalization in marketing, with the possible exception of the Bass diffusion model.

The power law is a central organizing concept of our analysis. In the context of brands, this law states that the market share, $s(r)$, of the brand with r-th highest market share is $s(r) = A(a + r)^{-b}$, where A, a, and b are constants. A key empirical content of the power law is that the decrease in the market share between two successively-ranked brands becomes smaller as one progresses from higher-ranked to lower-ranked brands.[1] A

[1]The power law has been employed to describe many patterns observed in the human, physical and biological world. We note some of these later. In the context of market shares of brands, Chung and Cox (1994) have examined the power law for the number of hit albums recorded by music groups. Kalyanaram, Robinson and Urban (1995) consider prescription anti-ulcer drugs and certain consumer packaged goods, and examine a special case of the power law relationship between the order of market entry and market shares. Some remarks on the connection between this study and our analysis are presented later.

testable alternative to the power law, which we examine, is that the ratio of market shares for any two successively-ranked brands is constant. This implies the exponential form; that is, the relationship between market shares and ranks is $s(r) = Ge^{-gr}$, where G and g are constants.

In this and the next several paragraphs, we summarize our main empirical findings, each of which is discussed at greater length in a later section. The key findings, when each product category is considered separately from others, are as follows:

1. The power law holds very well in an absolute sense; the R^2 values are consistently over 0.90 and often over 0.95.

2. The power law fits the data significantly better than the exponential form.

3. The relative superiority of the power law over the exponential form is greater for those product categories that have lower values of *b*. This finding supports Mandelbrot's (1963) theoretical result that a power law with sufficiently large values of *b* approximates the exponential form.

4. The exponential form fits better for lower-ranking market shares than for higher-ranking market shares.

In an important study on market shares, Buzzell (1981) tested the exponential form, without explicitly comparing it to an alternative specification. His analysis was based on market shares of firms calculated in dollars. Our analysis, as will be seen, is based on market shares of individual brands calculated in terms of unit sales. Like Buzzell, we find that the exponential model fits the data reasonably well. However, our analysis suggests that the power law is a better description of data than the exponential form, not only because the former consistently yields higher values of R^2 than the latter, but for other reasons as well. For example, we find that the exponential form typically provides bad fits to data at the "top" end of market shares. That is, if one considers the subset of brands with the largest market shares, then the ratios of successive market shares are not constant, as predicted by the exponential form, but yield a decreasing series of values. Next, recall the theoretical result of Mandelbrot (1963) that a power law with large values of the coefficient *b* approximates the exponential form. Conversely, the exponential form does not approximate a power law when *b* is small. Our analysis suggests that the value of *b* is small for some product categories and large for others. For those cases in which the values of *b* are small, we find that the exponential form provides especially poor fits to the data. Put another way, only one of the two empirical generalizations consistently performs better for all of the data, and it is the power law. For these and other reasons described in the paper (for example, the consistency of our findings with a theoretical model presented later, and the relationship of this model to the literature on consumer choice), our analysis suggests that, for the purpose at hand, the power law is a better empirical generalization than the exponential form.

Our analysis does not support two other hypotheses on market shares. The first is the "rule of three and four" advanced by the Boston Consulting Group (1987). It predicts that the three largest market shares will be 40%, 30% and 20%. The second is the hypothesis of Kotler (1977) that the top three brands will respectively have 40%, 30% and 20% market shares. We do not discuss these two hypotheses any further.

We also present findings from some "derived data sets" which we create by combining and separating the raw data on market shares across product categories in particular ways. The primary motivation for examining these derived data sets is that they allow us to test some predictions of the explanatory model that we propose later. As will be seen, this model is based on Bose-Einstein statistics, which is related to the well-known Dirichlet-multinomial model of brand purchases. The three derived data sets, which are discussed in detail later in the paper, are created as follows:

A. The first derived data set is created by pooling the market shares across all product categories.

B. The second data set is created by considering those brands that occupy the highest rank in their respective product categories.

C. The third data set is created by taking the averages of the market shares of brands that hold a particular rank in their respective product categories, and then considering these averages across the ranks.

Now recall the four results concerning the power law stated earlier. We show that all of these four results hold for each of the three derived data sets described above. These regularities would have been difficult to anticipate without explicit analysis because there are no *a priori* reasons to expect that strong patterns exist among the market shares of brands across product categories. As will be seen, these results from the derived data sets contribute to a better theoretical understanding of the power law. To our knowledge, an analysis of such derived data sets is unavailable in the literature.

The notable accuracy of the fits provided by the power law has potential implications for marketing practice and research, and also for the study of economic issues related to market shares, such as government policies towards market dominance. We discuss some of these matters in later sections of the paper. It is also possible that analogues of the results that we obtain from derived data sets exist for some phenomena other than market shares, such as the distribution of city sizes and of individuals' income. An analysis of these is beyond the scope of the present paper.

Organization of the paper. Section II presents some preliminaries. Section III describes the data. Section IV summarizes the empirical findings. Section V discusses a theoretical model for the explanation of power laws for market shares. Section VI presents some potential implications of our findings. The last section contains concluding remarks.

II. SOME PRELIMINARIES

We repeatedly refer to the power law and exponential form in the rest of this paper. Hence, for brevity of exposition, we write them here as the following numbered expressions:

$$s(r) = A(a + r)^{-b}\text{; and} \tag{1}$$

$$s(r) = Ge^{-gr}, \tag{2}$$

where $r = 1, 2, \ldots,$ denotes the rank of the brand with the highest market share, second-highest market share, and so on; and $s(r)$ is the value of the r-th highest market share. Sometimes the power law is expressed as $s(r) = A'(a' + hr)^{-b}$, where A', a', h and b are constants. This expression is the same as (1), as can be seen by setting $A \equiv A'h^{-b}$ and $a \equiv a'/h$. Next, define the ratio of two successively-ranked brands as $f(r) \equiv s(r)/s(r + 1)$. For brevity, we refer to $f(r)$ as the "share ratio." We assume that $A > 0$, $a > -1$, and $b > 0$ for the power law, and that $G > 0$ and $g > 0$ for the exponential form. These inequalities ensure that $s(r) > 0$ and $f(r) > 1$, and they are satisfied by all of our estimates.

The share ratio for the exponential form is $f(r) = e^g$, and that for the power law is

$$f(r) = (1 + (a + r)^{-1})^{b}. \tag{3}$$

It follows that for all r, $f(r) > f(r + 1)$ for the power law, and $f(r) = f(r + 1)$ for the exponential form. As was noted earlier, this is the central distinction between the power law and the exponential form: the ratio of market shares of two successively-ranked brands becomes smaller for the power law, but remains constant for the exponential form, as one progresses to lower-ranked brands.

Nomenclature and special cases. There does not seem to be a convention in the literature regarding the use of phrases such as the power law, the Pareto law, and Zipf's law. Sometimes these three phrases are used interchangeably, and sometimes more than one of these three phrases refers to one or another special cases of (1). For later use, we state some special cases of the power law, and do not use any particular phrases to refer to them. If $a = 0$, then from (1),

$$s(r) = Ar^{-b}. \tag{4}$$

This is perhaps the most commonly used version of the power law. The theoretical models cited later in Sections V and VII refer to this version. A special case of (4) is $s(r) = Ar^{-1}$, obtained when $b = 1$, which has been used extensively in the study of city sizes. Another special case of (4) is $s(r) = Ar^{-1/2}$, obtained when $b = 1/2$. As discussed later, this special case has been used by Kalyanaram, Robinson and Urban (1995) to study pioneering advantage. If the data under consideration is ranked in some specified sense, then the power law is sometimes called the "ranked power law" or the "rank-frequency form of the power law." In the present paper, the data is ranked by the magnitudes of market shares.

III. DATA

We examine two sets of data. The first, made available by Nielsen Market Research, contains the market shares for 506 brands in 48 product categories of foods. These market shares are for a large urban market in the Southwestern US, for a period of 120 weeks from January 1993 to May 1995. The second data set, published by the Sporting Goods Association of America, is for 665 brands in 43 product categories of sporting goods, for the calendar year 1999, for the entire US. All market shares are in equivalent (quantity) units; for example, gallons of orange juice and pairs of shoes. Table 1 displays the number of brands in each product category for foods; the names of the product categories are withheld by those who have provided the data. The product category number for foods, displayed in the first column of Table 1, does not play any role in this paper; it merely refers to the numbering of the product categories in the original data. Table 2 lists the names of the product categories for sporting goods and the number of brands in each product category. In these tables, we have displayed the product categories in descending order of the number of brands within a category. That is, the product categories with larger number of brands are displayed in the upper rows of the table, and those with smaller number of brands are displayed in the lower rows. This presentation of product categories will be helpful later.

As is the case with most data sets for market shares of brands, these two data sets reflect the particular motivations and constraints of those who created them. The following four aspects seem noteworthy in the present context: (a) The data on food categories, collected at the store level for a smaller geographical area, is perhaps more accurate than that on sporting-goods categories. A limitation of the former data is the exclusion of certain types of stores from the Nielsen audits. (b) The construction of product categories is not based on explicit considerations of empirical research. For example, we do not have a random selection of possible food and sporting-goods categories. The selection of product categories appears arbitrary for foods, and is biased towards different types of shoes among sporting goods. We have used all of the data that is available to us. (c) The available data does not distinguish brand variations and SKUs. For example, Minute Maid frozen orange juice is sold in different variations, such as with or without pulp and calcium, and in various sizes. The data adds the unit sales (that is, gallons) across variations and sizes, and the total is presented as the sale of Minute Maid brand within the product category of frozen orange juice. (d) The data are restricted to brands with market shares no smaller than 1%. This possibly excludes some local and regional brands.

The preceding caveats regarding the data sets are partly ameliorated by our findings. Foods and sporting goods are quite different; for example, regarding consumers' reasons for buying or not buying particular goods or brands, and the ways in which producers try to sell their products. Also, the data set for foods is regional while that for sporting goods is national. Furthermore, these two data sets have been constructed by two different organizations under different procedures and with different objectives, without any coordination with each other. Notwithstanding these key differences, the findings from the two sets of data are very similar. This could be viewed as a partial sign that our results and conclusions are likely to be robust with regard to the unique characteristics or limitations of the data sets.

Estimation method.[2] The power law expression, (1), can be rewritten as $\ln s(r) = \ln A - b \ln(a + r)$, which, given a, can be estimated using ordinary least squares. However, since the value of a is not known, we use a nonlinear estimation procedure to estimate the parameters and their standard errors. For the exponential form, we linearly estimate the parameter g by rewriting (2) as $\ln s(r) = \ln G - gr$.[3]

In certain cases, the market shares in a product category are identical up to two decimal places. In these cases, we assign the same average rank to these tied data points. For instance, if the market shares are identical for the brands with the third- and fourth-highest ranks, then we assign a rank of 3.5 to both of these market shares.

For several product categories, the number of brands is very small. For example, as seen in the lower rows of Table 1, there are 6 or fewer brands in each of 16 out of 48 product categories of foods. The empirical results presented in this paper should be viewed in light of several consequences of this limitation of data. For example, one would expect that the parameter estimates based on too few data points will likely have large standard errors. This typically turns out to be the case, especially for some of the estimates of the power law. Another consequence is as follows. For some of our estimates of the power law, the R^2 increases monotonically with the value of b. For such cases, we restrict the range of b to be no larger than 10.[4] This restriction is more often, but not always, binding

---

[2]Ideally, one would like to pool data across categories and simultaneously estimate all of the parameters using a random-effects model. However, the data at hand are not a random sample of food or sporting-goods categories.
[3]For brevity, we present the values of the estimated parameters, and those of the R^2, to only two places after the decimal.
[4]Beyond this value of b, there is very little marginal improvement in the fit in each of the cases we have examined. Estimates using larger upper bounds for b are available upon request from the authors. Also, note that it is theoretically possible that the optimal value of the estimate of b is

for product categories with fewer brands. We will discuss later another possible consequence of the fact that some of the product categories have a very small number of brands.

IV. EMPIRICAL RESULTS

In this section, we present the empirical findings. Many of these findings, especially those based on derived data sets, could not have been anticipated on an *a priori* basis. In the next section, we discuss some relationships between our empirical findings and a theoretical model of the power law.

Table 1 displays the estimated parameters, and the corresponding R^2, for the power law and the exponential form, separately for each of the 48 food categories. Table 2 displays the corresponding results for the 43 sporting-goods categories. These results suggest that the power law holds well in an absolute sense. For example, for foods, the R^2 for the power law is larger than 0.95 for 37 out of a total of 48 product categories, and it is larger than 0.9 for 44 product categories. For sporting goods, the R^2 for the power law is larger than 0.95 for 41 out of a total of 43 product categories, and it is larger than 0.9 for all 43 product categories.

Parameter estimates (of a, g and unrestricted b) that are statistically significant at the 5% confidence level are shown in boldface in Tables 1 and 2. Error estimates are not applicable if the reported value of b is 10, given the restriction mentioned earlier. The overall picture in this regard is that, if a parameter estimate is not significant, then it is typically but not always the case that the corresponding product category has a small number of brands. This can be seen in two different ways in Tables 1 and 2, in which, as was noted earlier, the product categories with larger numbers of brands are displayed in the higher rows of a table, and the product categories with smaller numbers of brands are displayed among the lower rows of the table. First, note that, compared to food categories, many more of the estimates of b are significant among sporting-goods categories, which typically also have more brands within individual categories than food categories. Second, note that the cases of significant estimates of b are more concentrated among the upper rows of these tables than among the lower rows. The small number of brands for many product categories seems to have a separate consequence with regard to whether or not the restriction on the estimated value of b is binding. Tables 1 and 2 suggest that this restriction is more likely to be binding for product categories which have a smaller rather than a larger number of brands.

---

smaller than 10, even if the restriction under consideration is binding. This is because if $q(A, a, b)$ denotes the value of the R^2, as a function of (A, a, b), which are the parameters to be estimated, then it is possible that $\text{Max}_{A,a} : q(A, a, 10) < \text{Max}_{A,a,b} : q(A, a, b < 10)$. However, for all of the estimates presented in this paper for which this restriction is binding, it turns out that the optimal value of the estimate of b is 10.

Figure 1 shows how the R^2 for the power law and the exponential form differ at different values of b. The upper panel in this figure is for food categories. The estimated values of b for food product categories are taken from Table 1, and the product categories are reordered according to ascending values of b, which are then displayed on the horizontal axis of this panel. The vertical axis displays the corresponding value of R^2 for the power law as well as for the exponential form. The numbers displayed on the horizontal axis are simply the labels of the product categories, after the reordering just stated; these numbers are in themselves not relevant to what this figure shows. The lower panel of Figure 1 presents the corresponding results for sporting-goods categories. This figure shows that the power law does substantially better, in terms of the R^2, than the exponential form for lower values of b, and that the two R^2 are less distinguishable from one another at higher values of b. These findings support the theoretical conclusion of Mandelbrot (1963) that a power law with sufficiently large values of b approximates the exponential form.

We now set up the apparatus which we will use below to describe, graphically as well as mathematically, the derived data sets which we create by separating and combining the raw data across product categories in particular ways. All of the derived data sets are industry specific; that is, the data for foods are not combined with the data for sporting goods. Figure 2 is a useful graphical aid in understanding these data sets; the upper panel is for foods and the lower panel is for sporting goods. This figure displays the market shares for each rank across product categories. That is, the highest-ranked market shares across product categories form the vertical cluster at rank = 1, the second-highest-ranked shares form another vertical cluster at rank = 2, and so on. Note that the neighboring clusters in Figure 2 overlap in their vertical ranges, but their means are quite well separated, as will be seen later in Figure 5.

The following notation is used below to describe formally the derived data sets: $c = 1$ to C denotes the product categories; R_c denotes the number of brands in product category c; $r = 1$ to R_c are the ranks of brands in product category c, where R_c is the rank of the lowest-ranked brand in product category c; $s_c(r)$ is the market share of the brand with rank r in product category c; and $H \equiv \text{Max}_c\, R_c$ is the largest number of brands in any of the product categories. Thus, as is seen in the first two columns of Table 1, $C = 48$ for foods, and $H \equiv 27$ is the largest number of brands in any food category. Define the integer $\rho = 1$ to H. For any ρ, define $\psi(\rho)$ as the set of those values of c for which $\rho \leq R_c$. That is, the product category c is included in the set $\psi(\rho)$ if it has ρ or more brands. Define $\phi(\rho)$ as the number of elements in the set $\psi(\rho)$. Then the data in Figure 2 is described as follows. The vertical cluster at rank = 1 displays the numbers $\{s_c(1) \mid c \in \psi(1)\}$, the vertical cluster at rank = 2 displays the numbers $\{s_c(2) \mid c \in \psi(2)\}$, and so on.

The first derived data set is presented in Figure 3. This figure displays the market shares of brands pooled across all product categories. Put differently, the observations contained in all of the vertical clusters in Figure 2 are combined together, and then rearranged in descending order, to produce the data points in Figure 3. Formally, the market shares displayed in Figure 3 are the numbers $\{s_c(r) \mid r = 1 \text{ to } R_c, \text{ and } c = 1 \text{ to } C\}$, rearranged in descending order. The upper two panels are for foods; the left panel is for the power law and the right panel is for the exponential form. The lower two panels are for sporting goods. In addition to the data, each panel presents the parameter estimates (namely, a and b for the power law, and g for the exponential form), the associated R^2, and a line which describes the pattern predicted by the estimated parameters. We follow the same conventions for graphical presentations in later figures.

Next, recall that the second derived data set was created by selecting the market shares of brands that hold the highest rank in their respective product categories. Figure 4 presents the results for this derived data set. Thus, the market shares displayed in this figure are the numbers which form the vertical cluster at rank = 1 in Figure 2, ordered in accordance with their rank within this cluster. Formally, these numbers are $\{s_c(1) \mid c \in \psi(1)\}$, rearranged in descending order.

The third derived data set is presented in Figure 5. Here we analyze the averages of the market shares of brands that hold a particular rank in their respective product categories. That is, the market share displayed at rank = 1 in Figure 5 is the average of the vertical cluster at rank = 1 in Figure 2. Formally, the data points shown in Figure 5 are the numbers $\{\sum_{c\in\psi(\rho)} s_c(\rho)/\phi(\rho) \mid \rho = 1 \text{ to } H\}$, rearranged in descending order.

It is readily seen from Figures 3, 4 and 5 that, for the derived data sets respectively described in these figures, the power law holds very well in an absolute sense. All of the estimates (of a, g and unrestricted b) presented in these figures are significant at the 5% level. For the derived data sets in Figures 3 and 5, the power law performs better than the exponential form. For the derived data set in Figures 4, the R^2 for the exponential form is comparable to that for the power law. Some comments on this last result are presented in the next section.

Finally, recall our conclusion that the exponential form fits better for lower-ranking market shares than for higher-ranking market shares. Consider Figure 3 as an illustration; analogous observations hold for other results presented in this paper, including those for individual product categories. In particular, consider the lines in Figure 3 which describe the patterns predicted by the estimated parameters of the exponential form. These lines are in the two right panels of Figure 3; the upper panel is for foods and the lower panel is for sporting goods. The fits provided by these lines are markedly better for lower-ranking market shares than for higher-ranking market shares. There is no such visual asymmetry in the fits provided by the power law, which are presented in the corresponding left panels of Figure 3. One way to interpret such asymmetry is as follows. As seen in Figure 3, there are significantly fewer observations for higher-ranked market shares than for lower-ranked market shares. Also, the changes in successively-ranked market shares exhibit bigger jumps for higher-ranked market shares. The asymmetry just noted with regard to the exponential form may therefore indicate an additional strength of the power law which is not entirely captured by the respective R^2 of these two specifications.

V. EXPLANATIONS

A large number of human and social phenomena have been described using the power law or its special cases.[5] Among them are the frequencies of words in texts (Zipf 1949); the counts of distinct names appearing in lists of manufacturers and distributors (Zipf 1950); citations of articles in academic journals (Simon 1955); the incomes, industrial capacities, and sizes of firms (Simon and Bonini 1958, Stanley et al. 1995, Ijiri and Simon 1974, and Okuyama, Takayasu and Takayasu 1999); city populations (Carroll 1982 provides a comprehensive review; see Gabaix 1999 and references cited therein for more recent work); actor collaborations (Barabasi and Albert 1999); dividends in horse races (Park and Domany 2001); and the distribution of individual incomes (see Persky 1992 for a review).

Each of these phenomena entail "micro-level" issues, including choices made by entities (such as individuals, intermediaries, and firms), the histories and expectations of these entities, and interactions among various entities. Given how disparate these phenomena are, a natural presumption should be that the micro-level issues are quite different for different phenomena. As Simon (1955, p. 425) notes, "No one supposes that there is any connexion between horse-kicks suffered by soldiers in the German army and blood cells on a microscope slide other than that the same urn scheme provides a satisfactory abstract model of both phenomena."

An objective of this section is to present a theoretical explanation for the observed power law for market shares. We seek an explanation that is consistent with the research literature on brand purchases; that leads to testable predictions; and that offers reasonable and intuitive insights. We present a model, due to Hill (1974), which we believe meets these criteria.[6] In a later section, we remark briefly on another class of theoretical models of the power law, namely, those based on Gibrat's law.

At the heart of Hill's model is the hypothesis that, within a product category, consumers' purchases of different brands are represented by Bose-Einstein statistics. A qualitative implication of this allocation mechanism is that if a certain number of units have already

[5]Many phenomena in the physical and biological worlds have also been described by the power law or its special cases. These include the magnitudes of avalanches, forest fires and earthquakes (the well-known log-log Richter scale); the distributions of fragments' masses in impact fragmentation of solid bodies; the masses of asteroids; and the global distribution of minerals. See Gell–Mann (1994) for a perspective.
[6]For vividness, this abstract model uses the language of biology; for example, the distribution of species across genera. Given our context, we use the language of market shares, brands, product categories, and so on.

been allocated in a given manner across brands, then the probability of the next unit being allocated to a particular brand is proportional to the total number of units that have already been allocated to this brand. In the next subsection, we present this implication in some detail; in particular, we work with the Dirichlet-multinomial model, of which Bose-Einstein statistics is a special case. We do this for two reasons. First, our presentation is a natural way to understand the implication of interest here. Second, our presentation makes it easier to understand a possible generalization of Hill's model, which we discuss later. We then briefly summarize the central assumptions of Hill's model.[7] This is followed by a discussion of some of the predictions of this model. At the end of this section, we discuss some qualitative aspects of this model as well as a possible generalization.

An implication of the Dirichlet-multinomial model. Consider one product category in isolation from others. Assume that purchases are made in increments of one unit. Let j = 1 to N denote the units and let i = 1 to M denote the brands. These M brands are not ranked in any manner, and N is larger than M. For brevity, we use the following shorthand. If i-th brand is selected when the j-th unit of the product is bought, then we refer to this event as: the j-th unit is allocated to brand i. It is important to keep in mind that we use the preceding shorthand solely for expositional convenience; there is no presumption that anyone is actually "allocating" product units to consumers. Instead, autonomous consumers are allocating purchases to brands.

Let L_i denote the number of units allocated to brand i. Let

$$L = (L_1, ..., L_M), \text{ where } \sum L_i = N \tag{5}$$

denote a particular allocation of N units across M brands. For brevity, in (5) and in the rest of this paper, we suppress the range of the index i over sums and products; i ranges from 1 to M.

We next ensure that at least one unit is allocated to each brand. Let U_j denote a mapping from units to brands. If the j-th product unit is allocated to brand i, then $U_j = i$. Assume that

$$U_1 = 1,\ U_2 = 2, U_3 = 3, \ldots, U_M = M; \tag{6}$$

that is, the first unit is allocated to the first brand, the second unit is allocated to the second brand, and so on. Expression (6) ensures, without any loss of generality, that at least one unit is allocated to each brand. After these units are allocated, there are $n \equiv N - M > 0$ units that remain to be allocated. Define a vector ℓ, such that

[7] A complete derivation of the model would be redundant and is also unnecessary for our purposes; interested readers are referred to Hill (1974).

$$\ell_i \equiv L_i - 1,\ \ell \equiv (\ell_1,\ldots,\ell_M), \text{ and } \Sigma\ell_i = n. \tag{7}$$

Note from (6) that each brand already has one unit allocated to it. Hence, ℓ is a particular allocation of the remaining n units among the M brands. The usefulness of the vector ℓ is that it can be considered without the concern that one or more brands will have no units allocated to them.

Define the vector $p \equiv (p_1,\ldots,p_M)$, where $1 \geq p_i \geq 0$, and $\Sigma p_i = 1$. Suppose that p has a Dirichlet distribution with parameters $(\alpha_1,\ldots,\alpha_M)$. That is,

$$\Pr(p) = \frac{\Gamma(\Sigma\alpha_i)\Pi p_i^{\alpha_i - 1}}{\Pi\Gamma(\alpha_i)}. \tag{8}$$

Next, suppose that the distribution of ℓ is multinomial with parameters p. Then, the density of ℓ, given p, is

$$\Pr(\ell \mid p) = n!\,\Pi\frac{p_i^{\ell_i}}{\ell_i!}. \tag{9}$$

The unconditional distribution of ℓ is Dirichlet-multinomial :

$$\Pr(\ell) = \frac{n!\,\Gamma(\Sigma\alpha_i)\Pi\Gamma(\ell_i + \alpha_i)}{\Gamma(n + \Sigma\alpha_i)\Pi\{\Gamma(\alpha_i)\Gamma(\ell_i + 1)\}}. \tag{10}$$

If $\alpha_i = 1$ for $i = 1, \ldots, M$, then (10) yields Bose-Einstein statistics:

$$\Pr(L \mid M, N) = 1\Big/\binom{N-1}{M-1}. \tag{11}$$

Our objective in this subsection is to show that (11) implies that, if a certain number of units have already been allocated (purchased) in a given manner across brands, then the probability of the next unit being allocated to a particular brand is proportional to the total number of units that were already allocated to this brand. As a step towards this objective, we calculate the probability density $\Pr(p \mid U_1,\ldots,U_{M+y})$, where $n > y > 0$. This is the probability density of p, given that y units have already been allocated in addition to the initial allocation of M units described in (6), according to the vector $(U_1,\ldots,U_{M+y})$. Define the vector $x \equiv (x_1,\ldots,x_M)$, with $\Sigma x_i = y$. This vector represents a particular allocation of y units, in addition to the initial allocation of M units. Define another vector

$X \equiv (X_1, \ldots, X_M)$, where $X_i \equiv x_i + 1$, and $\Sigma X_i = M + y$. This vector represents a particular allocation of $M + y$ units, including the initial allocation (6). Note that vectors x and X respectively are analogues of vectors ℓ and L defined earlier, with the difference that now the allocation of a total of $M + y$ units is being considered. Also note that the vectors x and X are functions of y. Since the order in which different units are allocated to a particular brand does not affect p,

$$\Pr(p \mid U_1, \ldots, U_{M+y}) = \Pr(p \mid x). \tag{12}$$

By Bayes' Theorem,

$$\Pr(p \mid x) = \frac{\Pr(x \mid p)\Pr(p)}{\Pr(x)}. \tag{13}$$

Substitution of (8), (9), and (10) in (13), and the use of (12) yields

$$\Pr(p \mid U_1, \ldots, U_{M+y}) = \frac{\Gamma\left(\Sigma x_i + \alpha_i\right)\Pi p_i^{x_i+\alpha_i-1}}{\Pi\Gamma(x_i + \alpha_i)}. \tag{14}$$

Comparison of (14) with (8) shows that the distribution described in (14) is Dirichlet with parameters $(x_1 + \alpha_1, \ldots, x_M + \alpha_M)$. If the operator E denotes the expected value, then a standard property of the Dirichlet distribution is that

$$E(p_m \mid U_1, \ldots, U_{M+y}) = \frac{x_m + \alpha_m}{y + \Sigma\alpha_i}, \text{ for } m = 1, \ldots, M. \tag{15}$$

We finally consider the object of our interest: $\Pr(U_{M+y+1} = m \mid U_1, \ldots, U_{M+y})$. This is the probability that the next unit will be allocated to brand m, given that a certain number of units have already been allocated to various brands in a particular manner. The incremental unit under consideration is allocated to brand m with probability p_m, which is a random variable. Hence,

$$\Pr(U_{M+y+1} = m \mid U_1, \ldots, U_{M+y}) = E(p_m \mid U_1, \ldots, U_{M+y}). \tag{16}$$

Now recall (15) and that $X_i = x_i + 1$. It follows from (16) that

$$\Pr(U_{M+y+1} = m \mid U_1, \ldots, U_{M+y}) = \frac{X_m + (\alpha_m - 1)}{M + y + \Sigma(\alpha_i - 1)}. \tag{17}$$

Since, for Bose-Einstein statistics, α_i = 1 for all *i*, the probability that the next unit is allocated to brand m is proportional to X_m, which is the number of units already allocated to brand m.

Hill's model. The central assumption of this model is as follows: within a product category, units are allocated across brands according to Bose-Einstein statistics, given the total number of units to be allocated to various brands, and the number of brands, and given that the number of units is much larger than the number of brands. A classical alternative to Bose-Einstein statistics is Maxwell-Boltzmann statistics. Hill's analysis suggests that, within his framework, Maxwell-Boltzmann statistics will yield outcomes that are markedly different from those of the power law.[8]

Hill's model considers multiple product categories simultaneously, and allows the total number of brands in each category to be a random variable, thus accommodating the introduction and withdrawal of brands from the market for various product categories. Let C denote the number of product categories, where C is large. Let M_c denote the number of brands in category c, where $C \geq c \geq 1$. Let N_c denote the total number of units sold in category c. The main assumptions of the model with regard to product categories are as follows. The allocations within a product category are independent across product categories. The number of brands within a category, M_c, is random. The ratio M_c/N_c is independent across product categories. Next, define the distribution $\Pr(M_c / N_c \leq w \mid N_c)$. For any c, let $F(w)$ denote the corresponding limiting distribution as $N_c \to \infty$. A sufficient condition that F satisfies is that, for some positive constants γ and V, $F(w)/Vw^{\gamma} \to 1$ as $w \to 0$. There is a large family of distributions which satisfies the preceding condition. For example, γ = 1 includes the uniform distribution, and γ = 1/2 includes the arcsin distribution.

Predictions. An important aspect of Hill's model for the study of market shares is that it yields some testable predictions. One prediction is that the power law arises (in an approximate sense) when the market shares are pooled across product categories. As shown in Figure 3, our analysis supports this prediction. A result of Hill leads to another prediction, namely, that the power law arises among those markets shares which are the r-th largest in their respective product categories, provided that the value of r is low. We have tested this hypothesis for the highest ranked market shares, for which r = 1 (by definition), across product categories. As was noted earlier, the power law does quite well in an absolute sense. However, it does not do better than the exponential form. It is possible that this is because several of our product categories have very few brands, and that such limitations of data are inconsistent with some of the asymptotic arguments

[8] See Feller (1968, pp. 38–43) for a discussion of Bose-Einstein statistics and Maxwell-Boltzmann statistics.

required to establish the predictions under consideration. We plan to investigate this matter further in future research.

We have presented a pattern in Figure 5 which is not predicted in Hill (1974). Recall that, in this figure, we have: (i) calculated the mean of the k-th largest market shares across product categories, and (ii) arranged the means for various values of k in descending order. Our analysis suggests that the power law is a reasonable description of this data.

Additional observations. Hill's model has several strengths as a potential starting point for the study of market shares. The model yields some testable predictions, and our analysis of the derived data sets supports some of these. The model can be motivated in terms of consumers' purchases of different brands. The model can also be motivated in terms of multiple product categories, the introduction of new brands in various product categories, and the withdrawal of existing brands. At the same time, such motivation must be understood with caution because the model is a reduced-form model. What we mean by the phrase "reduced-form" here is as follows. The primary sources of the emergence of power laws in the model are particular statistical relationships, and the model does not explicitly deal with micro-level issues of the kind noted at the beginning of this section.[9]

Finally, recall that a central assumption of Hill's model is that Bose-Einstein statistics describe the purchases of unit sales of different brands within a product category. A conjecture of Hill is that the power law results will hold even if Bose-Einstein statistics are replaced by the Dirichlet-multinomial model. To our knowledge, this conjecture has not been proven or refuted. If true, this greatly increases the scope of Hill's model, because (10) represents a wide class of distributions. Also, as discussed earlier, the Dirichlet-multinomial model is a widely used and tested model in marketing research. This suggests a natural connection between the extensive literature on brand choices and the analysis presented in this paper.

VI. IMPLICATIONS

One might well ask if the present results are of any practical use. Suppose one were to accept that the market shares of brands follow a power law, then how might a manager use this result? We discuss these questions and others in this section, with the caveat that an explicit analysis of these implications requires data on and dynamic analysis of when and how market shares change.

[9] An analogous observation applies to a class of models of the power law based on Gibrat's law. We briefly discuss this class later.

From a managerial standpoint, the first implication of our analysis is that a brand manager has to view the stakes in her market share in discontinuous or "quantum" steps. A typical viewpoint is that incremental efforts produce incremental returns. The present paper appears to suggest otherwise. Incremental increases in effort are likely to produce at best transitory changes in market share that will average out to little more than "noise" over a relatively long period of time. This is consistent with what we know from the analysis of panel data, where we find large variations in weekly market shares, often contemporaneous with the promotional activities of firms, but little change over a year. To a frustrated brand manager, this may appear to be an indication of churning and of low brand loyalty in her particular market segment. In contrast, our results suggest that this is likely to be the norm across markets. Further, if there are significant changes, then a larger brand may face bigger changes in its market share. Thus, managing a leading brand within a given product category is likely to be a far riskier endeavor than managing a small brand. Whether due to luck or planning, market leaders have to play for big stakes, and this ought to be reflected both in the resources devoted to brand management and in the types of people employed to manage large brands.

Next, there is substantial evidence that market shares are generally stable in mature markets.[10] This suggests that the times when firms are more likely to be able to change their fortunes are before a market matures, or when a mature market becomes unstable for reasons such as rapid technological changes or marked reconfigurations in the rules of international trade. It may also be the case that, at least among the market leaders in a product category, changes in the ranking require large efforts and/or mistakes by challengers and/or by incumbents, and that incremental efforts and events may be less important to a brand than one might otherwise think.

Another implication of our results is that there are "natural" values around which brands can expect the market shares to be distributed. This means that the appearance of large deviations from the norm indicates either exceptionally good or bad performance, or instability in the market. In this sense, knowing the power law for a product category can provide a benchmark for brand performance. Such a benchmark can be useful to managers assessing the performance of a brand and also to financial analysts whose assessment of brand performance affects the stock prices of firms.

Our results also provide some insights into the sources of pioneering advantage. Kalyanaram et al. (1995) examine the relationship between firms' order of entry into the market and their respective market shares. They study prescription anti-ulcer drugs and certain consumer packaged goods. They estimate a special case of the power law, namely, expression (4) with $b = 1/2$, where r is the order of entry into the market. They find that the brands introduced earlier have on average a higher market share, represented

[10]See Lal and Padhmanabhan (1995), Dekimpe and Hanssens (1995), and Ehrenberg (1988), among others.

by the relationship just stated.[11] Our analysis finds that the power law is a much more widespread phenomenon; it holds for a much larger set of product categories, regardless of pioneering advantage; that is, regardless of considerations of the order of market entry. This suggests the need to separate the nature of pioneering advantage from the more general pattern of market shares which we have examined.

VII. CONCLUSIONS

The patterns of market shares described in this paper are tentative because, among other things, we have examined a total of only 91 product categories in two (relatively broad) markets. It is an open question whether such patterns exist in more comprehensive and better data sets. With this caveat, we conclude the paper with some speculative remarks.

At the phenomenological level, firms are concerned, often on an ongoing basis, with variables such as targeting and product positioning, product quality and brand equity, pricing and promotions, advertising expenditures and distribution intensity. A substantial body of research literature has attempted to understand the market shares of brands, using variables such as those just noted;[12] for brevity of exposition, we will refer to these approaches as "causational." This research also includes, at various levels of explicitness, considerations such as the histories of firms, the strategic interplay among firms, the behavior of consumers and intermediaries, the dynamics of product growth and innovation, and different kinds of uncertainties and expectations. Much of value has been learned from this literature, and will continue to be learned from its future developments.

We believe that our analysis complements the above literature. For instance, the high accuracy of power laws in describing the market shares of brands suggests that one should model the effects of marketing instruments through their effect on the parameters of a power function. A broader perspective is this. A macroscopic study such as ours deals with where brands end up in terms of market shares, and not with how they get there. A focus of causational studies is to understand the relationship between the firms' market shares and their efforts and environments. At some stage in the future, we would expect these two approaches to converge.

We have not emphasized other theoretical models of the power law because we believe that the line of theoretical reasoning presented in Hill's model is a particularly useful starting point in a marketing context. As discussed earlier in the paper, this model has the

[11] Kalyanaram et al. (1995) note that their relationship does not always hold for every specific data series they examine, but it does for the average market shares ordered by rank.

[12] A full discussion of this literature will take us far afield. Among the studies are Guadagni and Little (1983), Jedidi, Mela and Gupta (1999), Lancaster (1990), McFadden (1986), Robinson and Fornell (1985), and Urban, Carter, Gaskin and Mucha (1986).

advantage of generating testable propositions that are supported by the data, and its key assumption (namely, Bose-Einstein statistics) mirrors a basic paradigm of consumer behavior, namely, the Dirichlet-multinomial model of preferences. A well-known class of models of the power law is derived, directly or indirectly, from Gibrat's law. A discussion of the large and venerable literature on this class of models would be out of place here. In the context of market shares, Gibrat's law is: $S^i(t+1) = S^i(t)Z$, where $S^i(t)$ represents the unit sales of brand i in period t, and the random variable Z does not depend on i or on the unit sales of any brand.[13] This law posits that period-to-period changes in the sales of a brand follow a process in which the probability of any specified percentage change is independent of the brand's present size. In other words, the sales of each brand have the same probability of increasing or decreasing by 5 percent, 10 percent, or any other fixed amount, regardless of current sales. At least on the surface, this appears to be not entirely inconsistent with how firms might make resource allocation decisions.

We conclude with a remark on the role of chance. On the face of it, there is not much in common between the market for soft drinks and the market for handguns in the United States, including the sizes and characteristics of the firms in each of the two markets. We nevertheless find that the power law describes the patterns of market shares in these markets as well as in a number of other disparate product markets considered in this paper. The power law also holds in a number of different ways across product markets, as was shown by our analysis of the derived data sets. This suggests that chance plays a deep role. Kendall (1961, p. 12) has the following to say on the role of chance: "In fact, not only can choice mimic chance, but chance can mimic choice. Consider, for example, a number of persons with equal stakes playing at a fair zero-sum game (A zero-sum game is one in which the losses of some players pass to others, so that no money is lost to the system).[14] Over a period of play, the distribution of their holdings will tend to an unequal pattern of the Pareto type; some people will be reduced to small or zero stakes and a few will accumulate large reserves. And if the game goes on long enough, ultimately all the

---

[13] Additional assumptions are added, sometimes depending on the context in which Gibrat's law is being employed. Among the assumptions are those concerning: the support of *Z*; bounds on $S^i(t)$; births and deaths of entities (for example, Simon and Bonini 1958; for market shares, this aspect could be interpreted as the introduction of new brands and withdrawals of existing brands); intertemporal correlations among growth rates (Ijiri and Simon 1974); and intertemporal variability in the distribution of *Z*. Separately, recall a remark made earlier regarding the motivations associated with models of the power law. Gibrat's law can be motivated using analogies of intertemporal expansions and contractions of brands. Some models of the power law belonging to this class analyze some micro-level issues, and such analyses are valuable for their respective contexts of application. The power law, however, is derived from the statistical specification of Gibrat's law. An example of such analysis of micro-level issues, in the context of city populations, is Gabaix (1999). His focus is on the special case of (6) where $b = 1$.

[14] Evolution of market shares are of course zero-sum changes; firms' gains and losses must add up to zero. (This footnote is not part of the quotation.)

money will be concentrated in a few hands. Thus chance, which might be expected to level things out, will act in a very inegalitarian way and produce a distribution with a very purposive-looking outcome. It is not Providence but Chance which is on the side of the big battalions."

TABLE 1

Estimates for each product category: Foods

Product Category	No. of brands	Exponential		Power Law		
		g	R^2	a	b	R^2
8	27	**0.06**	0.84	-0.17	**0.59**	0.97
39	22	**0.11**	0.91	1.01	**1.10**	0.97
9	21	**0.12**	0.84	**-0.65**	**0.75**	0.91
30	21	**0.12**	0.97	16.92	**3.35**	0.98
14	20	**0.15**	0.91	**1.22**	**1.47**	0.99
31	20	**0.15**	0.95	**4.78**	**2.03**	0.99
40	20	**0.12**	0.92	2.35	**1.31**	0.95
12	19	**0.13**	0.89	0.40	**1.06**	0.99
42	17	**0.18**	0.96	**3.66**	**2.00**	0.99
2	16	**0.18**	0.98	17.04	**4.52**	0.99
36	16	**0.20**	0.97	9.56	**3.41**	0.98
4	13	**0.25**	0.90	0.87	**1.63**	0.98
43	13	**0.23**	0.96	4.40	**2.38**	0.98
1	12	**0.20**	0.90	0.75	**1.23**	0.98
11	12	**0.24**	0.90	0.31	**1.33**	0.98
32	12	**0.29**	0.99	**28.58**	10.00	0.99
10	11	**0.30**	0.94	**28.08**	10.00	0.94
17	11	**0.30**	0.99	**27.33**	10.00	0.99
23	11	**0.36**	0.94	1.98	**2.57**	0.98
45	11	**0.28**	0.90	0.89	**1.62**	0.94
27	10	**0.38**	0.83	**-0.61**	**1.28**	0.97
34	10	**0.40**	0.96	**19.85**	10.00	0.96
35	10	**0.41**	0.97	**19.29**	10.00	0.97
7	9	**0.35**	0.93	4.57	3.20	0.93

Product Category	No. of brands	Exponential		Power Law		
		g	R^2	a	b	R^2
6	8	**0.55**	0.92	2.92	3.82	0.95
20	8	**0.50**	0.97	8.29	6.31	0.98
38	8	**0.48**	0.95	**16.85**	10.00	0.95
46	8	**0.49**	0.98	15.00	9.39	0.98
16	7	**0.32**	0.92	**27.89**	10.00	0.92
19	7	**0.51**	0.98	**16.07**	10.00	0.98
26	7	**0.61**	0.85	3.79	4.60	0.87
29	7	**0.68**	0.96	**10.97**	10.00	0.96
3	6	**0.55**	0.96	**15.00**	10.00	0.96
5	6	**0.86**	0.95	3.41	5.72	0.97
15	6	**0.78**	0.96	1.93	4.00	0.99
18	6	**0.55**	0.99	12.36	8.67	0.99
33	6	**0.71**	0.91	-0.16	**1.94**	0.99
37	6	**0.81**	0.86	**9.45**	10.00	0.86
25	5	**0.92**	0.94	0.22	2.57	0.97
28	5	**0.85**	1.00	**8.86**	10.00	1.00
41	5	0.99	0.69	**-0.99**	**0.75**	1.00
44	5	**0.78**	0.89	0.02	1.99	0.91
47	5	**0.74**	0.92	**10.71**	10.00	0.92
48	5	**1.01**	0.97	**7.26**	10.00	0.97
13	4	0.59	0.69	15.02	10.00	0.68
21	4	0.55	0.83	-0.91	0.48	0.96
22	4	**1.19**	0.99	**6.12**	10.00	0.99
24	4	0.77	0.83	10.94	10.00	0.82

TABLE 2

Estimates for each product category: Sporting Goods

Product Category	No. of brands	Exponential g	Exponential R^2	Power Law a	Power Law b	Power Law R^2
Sports sandals	37	**0.07**	0.92	**3.37**	**1.29**	0.99
Boat shoes	37	**0.07**	0.94	**3.95**	**1.32**	0.99
Backpacks	36	**0.09**	0.92	**3.29**	**1.62**	0.96
Walking shoes	34	**0.09**	0.86	**0.61**	**1.22**	0.99
Hiking boots	34	**0.08**	0.89	**1.42**	**1.10**	0.98
Sneakers (gym shoes)	31	**0.10**	0.85	**0.48**	**1.21**	0.99
Fitness shoes	30	**0.12**	0.88	0.38	**1.37**	0.99
Tennis shoes	27	**0.12**	0.86	0.28	**1.24**	0.99
Scooters	26	**0.12**	0.75	**-0.88**	**0.85**	0.97
Aerobic shoes	24	**0.15**	0.83	-0.02	**1.30**	0.96
Golf clubs	19	**0.14**	0.94	**2.24**	**1.49**	1.00
Soccer balls	18	**0.17**	0.95	**4.78**	**2.25**	0.98
Cross training shoes	18	**0.23**	0.86	0.28	**1.76**	0.99
Sleeping bags	16	**0.18**	0.72	**-0.97**	**0.63**	0.95
Golf club sets	16	**0.14**	0.98	21.73	4.06	0.98
Jogging shoes	16	**0.28**	0.91	2.67	**2.81**	0.97
Golf bags	15	**0.17**	0.99	**12.53**	**3.25**	0.99
Skateboarding shoes	15	**0.22**	0.79	**-0.56**	**1.09**	0.98
Hunting boots	15	**0.13**	0.87	0.08	**0.77**	0.97
Soccer shoes	13	**0.32**	0.88	0.05	**1.73**	0.98
Fishing reels	13	**0.24**	0.97	16.67	5.60	0.98
Basketball shoes	12	**0.42**	0.89	0.43	**2.33**	0.97

Product Category	No. of brands	Exponential g	Exponential R^2	Power Law a	Power Law b	Power Law R^2
Baseball shoes	12	**0.35**	0.87	0.08	**1.79**	0.97
Tents	11	**0.31**	0.93	0.46	**1.63**	0.98
Rifles	9	**0.31**	0.99	**27.67**	10.00	0.98
Water sport shoes	9	**0.14**	0.95	**64.53**	10.00	0.95
Golf shoes	9	**0.42**	0.94	0.77	**2.08**	0.99
Treadmills	8	**0.43**	0.84	**-0.68**	**1.07**	0.99
Ice hockey sticks	8	**0.27**	0.98	**32.64**	10.00	0.98
Handguns	8	**0.30**	0.97	**29.03**	10.00	0.97
Cycling shoes	8	**0.23**	1.00	**39.39**	10.00	1.00
Bowling shoes	8	**0.54**	0.90	0.46	2.03	0.94
Billiard pool sticks	8	**0.19**	0.95	-0.07	**0.59**	1.00
Baseball softball bats	8	**0.37**	0.94	**22.69**	10.00	0.93
Baseball mitts	8	**0.50**	0.95	**15.58**	10.00	0.95
Tennis racquets	7	**0.53**	0.91	-0.02	**1.73**	0.98
Track shoes	7	**0.50**	0.95	0.09	**1.63**	1.00
Cheerleading shoes	7	**0.28**	0.66	**-0.98**	**0.37**	0.99
Basketballs	7	**0.65**	0.95	11.20	9.80	0.96
Volleyballs	6	**0.25**	0.97	0.48	0.84	1.00
Footballs	6	**0.62**	0.91	0.11	1.89	0.97
Football shoes	5	**0.95**	0.95	1.97	4.42	0.96
Air pistols	4	**1.29**	0.91	**32.64**	10.00	0.98

FIGURE 1

R^2 values for each product category: power law and exponential form

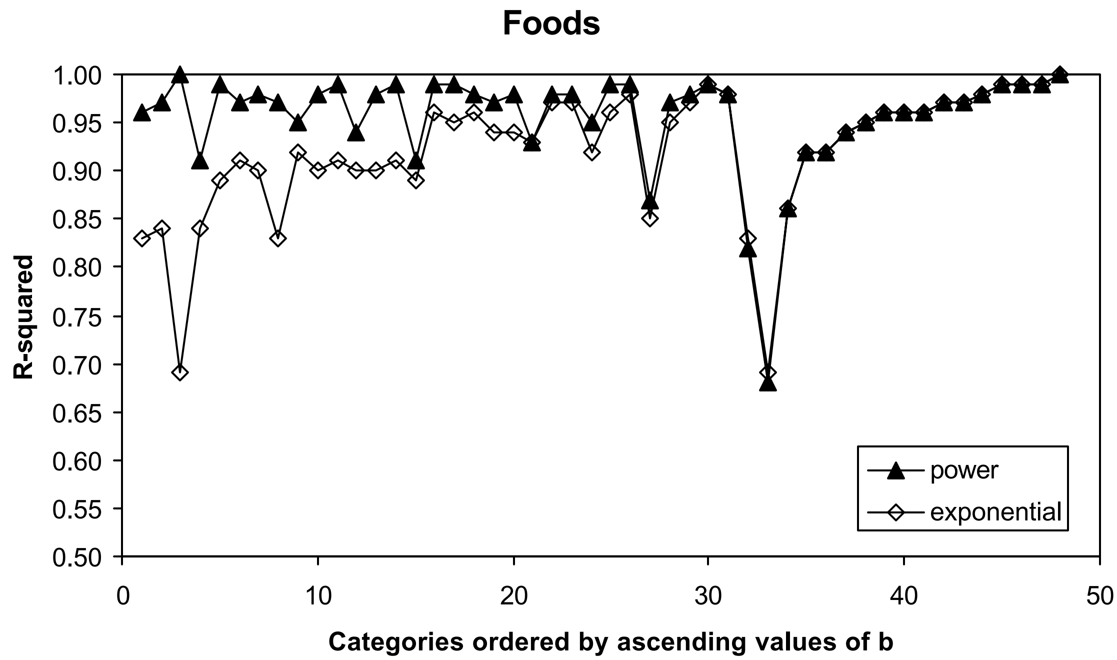

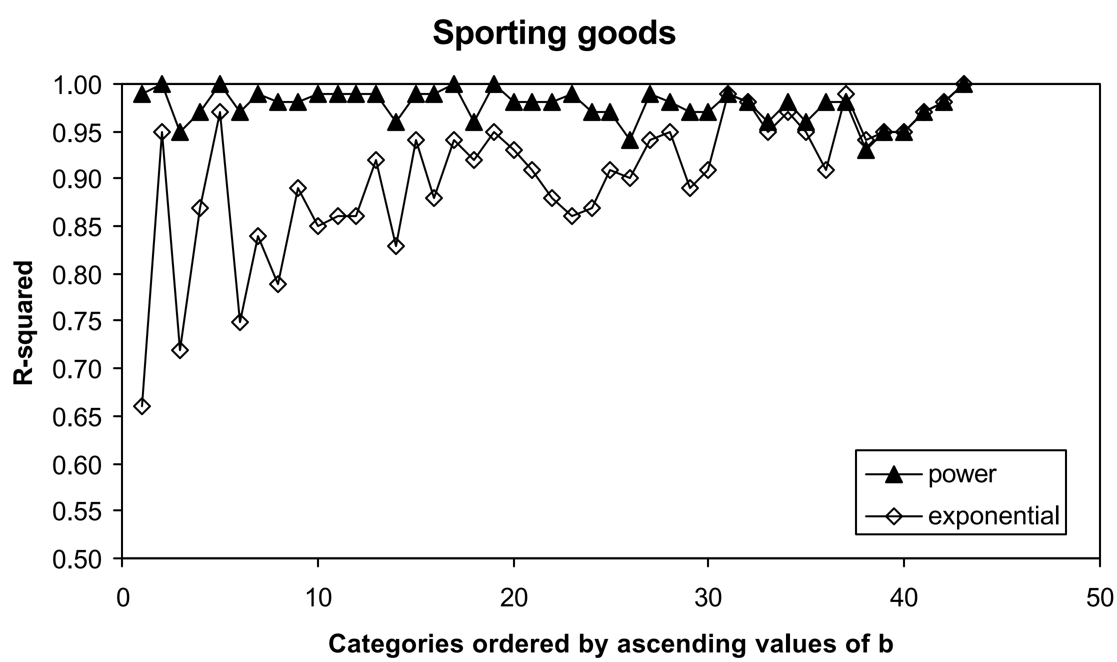

FIGURE 2

Market shares by rank, across product categories

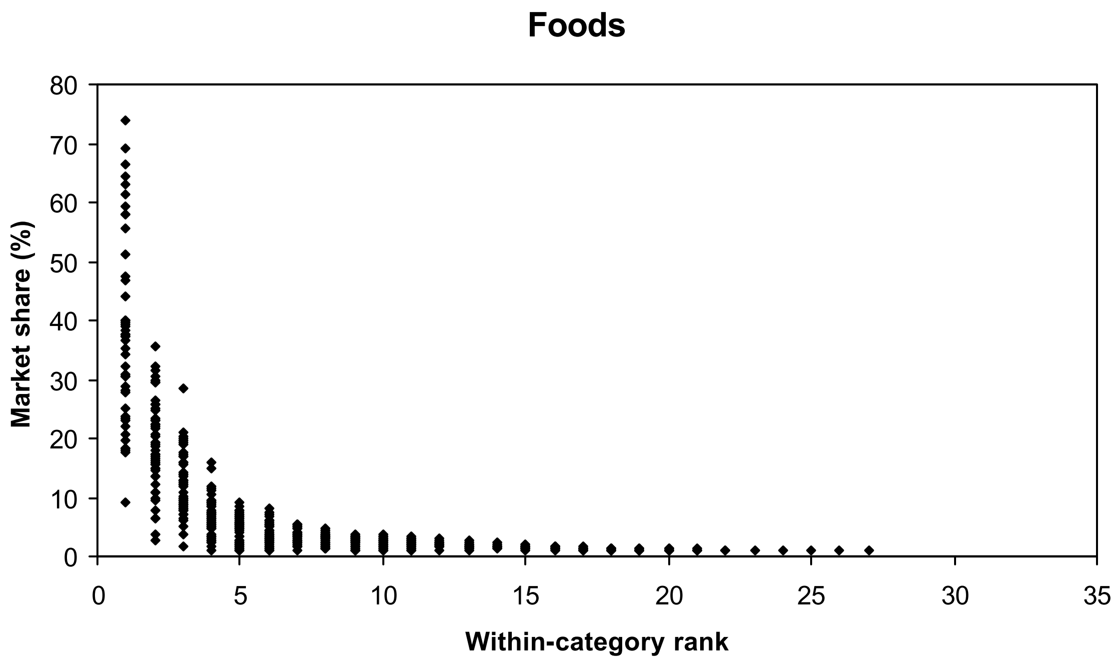

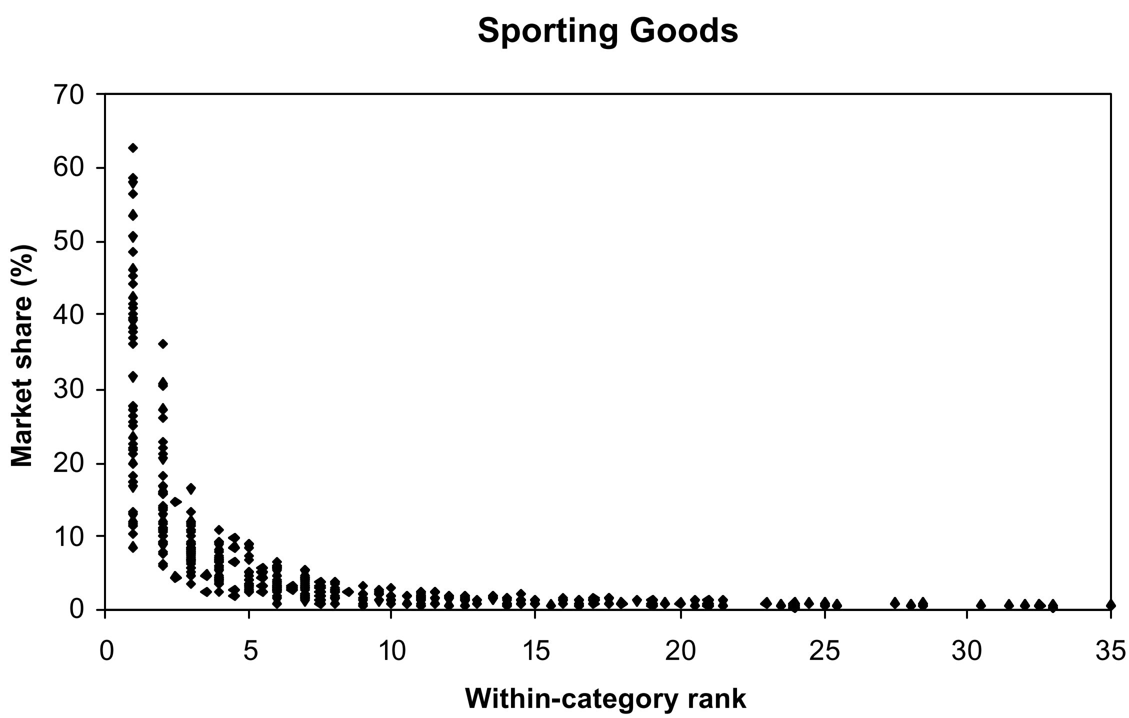

FIGURE 3

Market share versus rank across 506 food brands and 665 sporting-goods brands

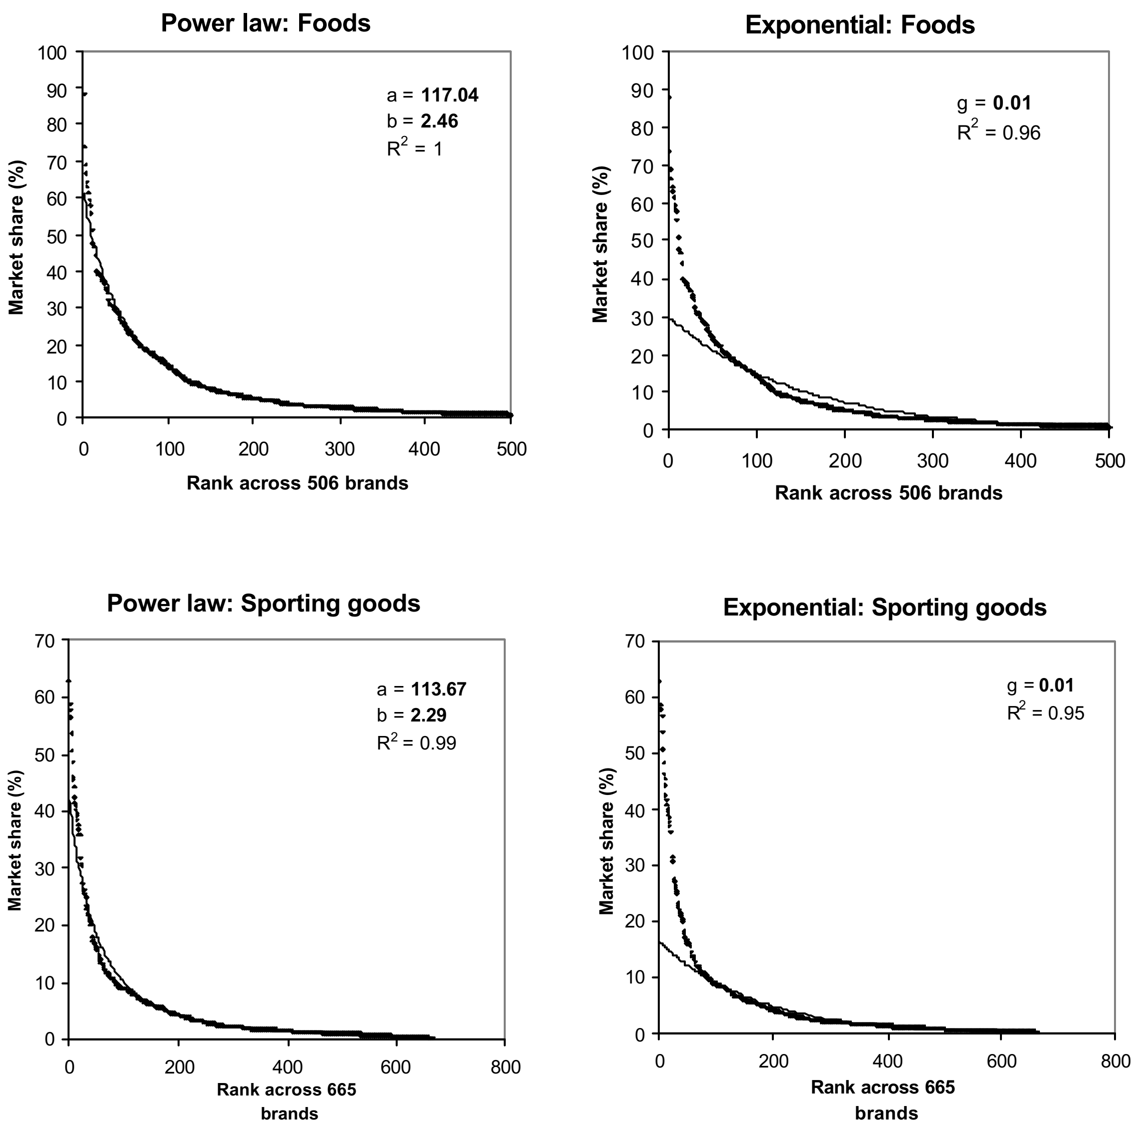

FIGURE 4

Highest-ranked market shares, across product categories

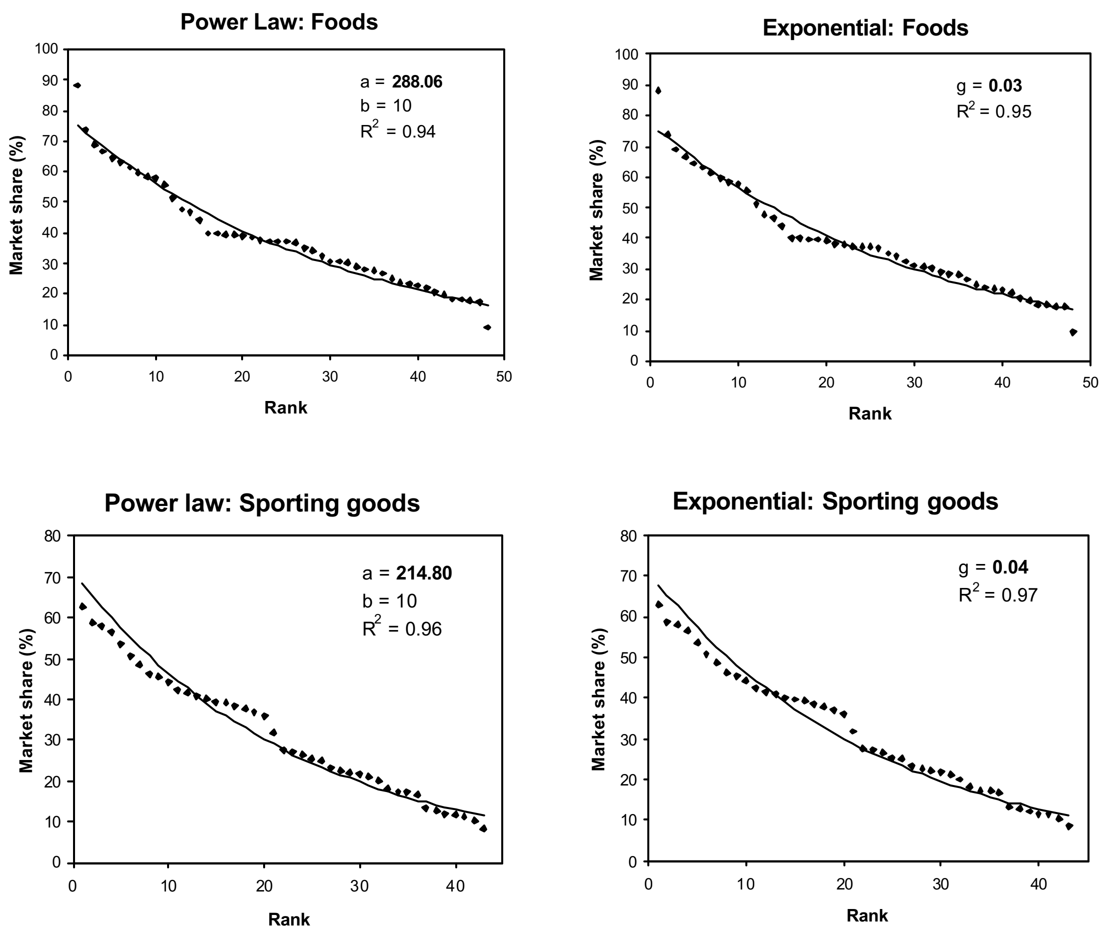

FIGURE 5

Average market shares with same rank, across product categories

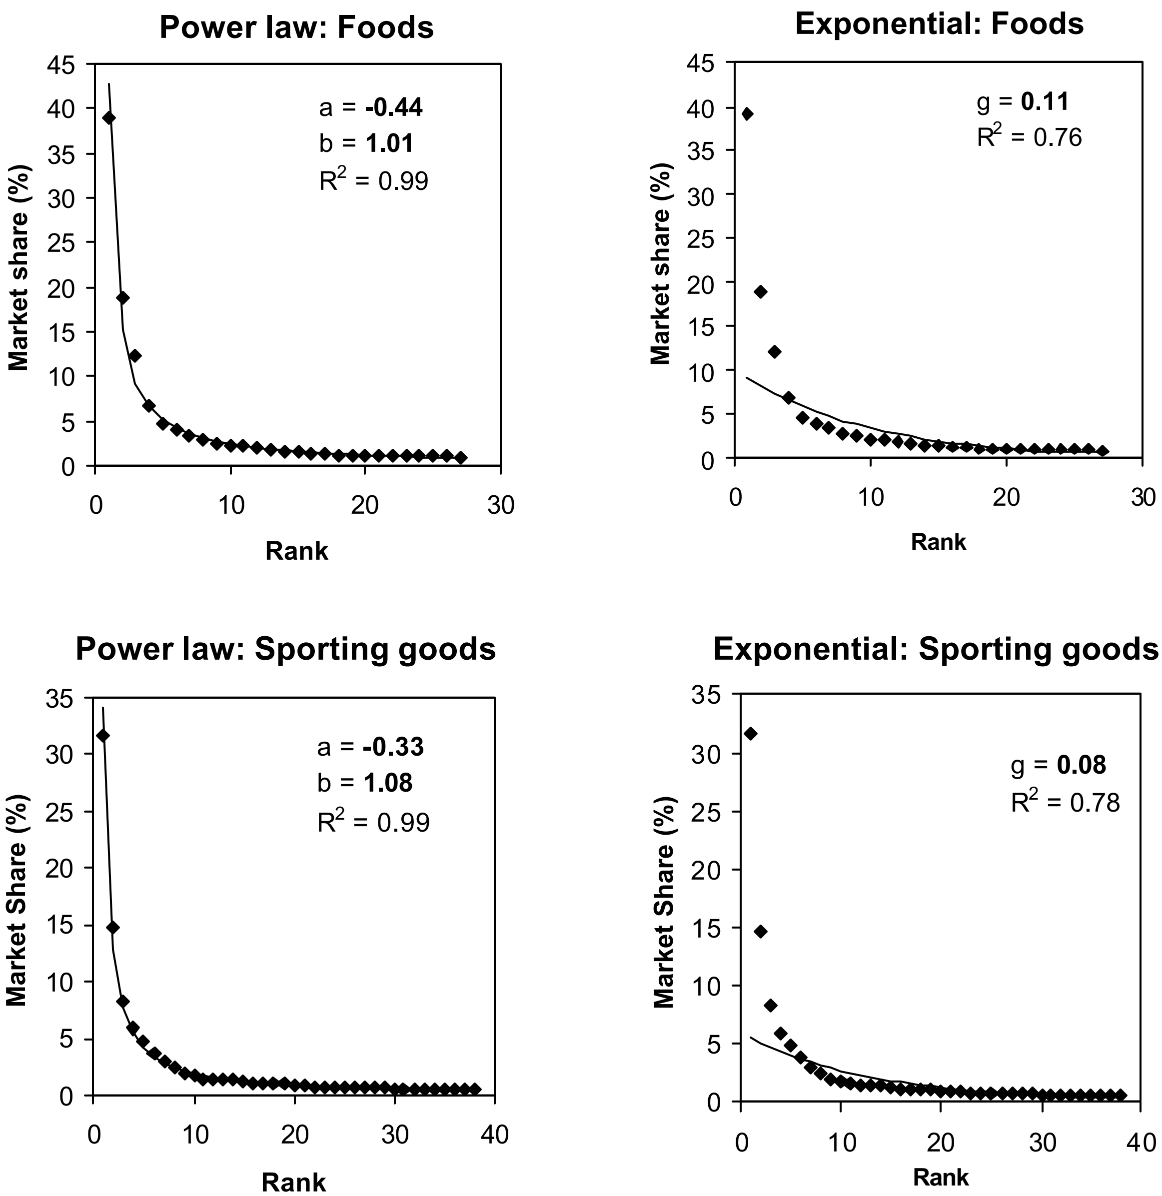

REFERENCES

Barabasi, Albert-Laszlo and Reka Albert (1999), "Emergence of Scaling in Random Networks," *Science*, 286 (5439), 509-512.

Bass, Frank M. (1995), "Empirical Generalizations and Marketing: A Personal View," *Marketing Science,* 14 (3) Part 2, G6-G19.

Bass, Frank M. (1969), "A New Product Growth Model for Consumer Durables," *Management Science,* 15 (1), 215-227.

Boston Consulting Group (1987) "The Rule of Three and Four," *Perspectives* 187, Boston MA.

Buzzell, Robert D. (1981), "Are There 'Natural' Market Structures?" *Journal of Marketing*, 45 (1), 42-52.

Carroll, G. (1982), "National City-Size Distribution: What Do We Know After 67 Years of Research?" *Progress in Human Geography,* 6, 1-43.

Chatfield, Chris and G. J. Goodhardt (1975), "Results Concerning Brand Choice," *Journal of Marketing Research,* 12 (1)*,* 110-113.

Chung, Kee H. and Raymond A. K. Cox (1994), "A Stochastic Model of Superstardom: An Application of the Yule Distribution," *Review of Economics and Statistics,* 76 (4), 771-775.

Dekimpe, Marnik G. and Dominique M. Hanssens (1995), "Empirical Generalizations About Market Evolution and Stationarity," *Marketing Science*, 14 (3), Part 2, G109-G121.

Ehrenberg, Andrew S. C (1995), "Empirical Generalizations, Theory and Methods," *Marketing Science,* 14 (3) Part 2, G20-G28.

Ehrenberg, Andrew S. C. (1988), *Repeat Buying: Facts, Theory and Data*, 2nd Edition, New York NY: Oxford University Press.

Ehrenberg, Andrew S. C (1982), "Lawlike Relationships," Research in Marketing Series, No. 82/1, London Business School.

Feller, William (1968), *An Introduction to Probability Theory and its Applications,* Volume I, 3rd edition, New York NY: Wiley.

Gabaix, Xavier (1999), "Zipf's Law for Cities," *Quarterly Journal of Economics*, 114 (3), 738-767.

Gell-Mann, Murray (1994), *The Quark and the Jaguar*, New York NY: Freeman.

Goodhardt, G. J., Andrew S. C. Ehrenberg and Chris Chatfield (1984), "The Dirichlet: A Comprehensive Model of Buying Behavior," *Journal of the Royal Statistical Society* A, 147 (5), 621-655.

Guadagni, Peter M. and John D. C. Little (1983), "Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science,* 2 (3), 203-238.

Hill, Bruce M. (1974), "The Rank-Frequency Form of Zipf's Law," *Journal of the American Statistical Association,* 69 (348), 1017-1026.

Ijiri, Yuji and Herbert A. Simon (1974), "Interpretations of Departures from the Pareto Curve Firm-Size Distribution," *The Journal of Political Economy,* 82 (2), Part I, 315-331.

Jedidi, Kamel, Carl F. Mela and Sunil Gupta (1999), "Managing Advertising and Promotion for Long-Run Profitability," *Marketing Science,* 18 (1), 1-22.

Kalyanaram, Gurumurthy, William T. Robinson and Glen L. Urban (1995), "Order of Market Entry: Established Empirical Generalizations, Emerging Empirical Generalizations, and Future Research," *Marketing Science,* 14 (3), Part 2, G212-G221.

Kendall, M. G. (1961), "Presidential Address," *Journal of the Royal Statistical Society: Series A (General)*, 124 (1), 1-16.

Kotler, Philip (1977), *Marketing Management: Analysis, Planning, Implementation and Control,* Englewood Cliffs, NJ: Prentice Hall.

Lal, Rajiv and V. Padmanabhan (1995), "Competitive Response and Equilibria," *Marketing Science*, 14 (3), Part 2, G101-G108.

Lancaster, Kelvin (1990), "The Economics of Product Variety: A Survey," *Marketing Science,* (9) 3, 189-206.

Mandelbrot, Benoit (1963), "New Methods in Statistical Economics," *The Journal of Political Economy,* 71 (5), 421-440.

McFadden, Daniel (1986), "The Choice Theory Approach to Market Research," *Marketing Science,* 5 (4) , 275-297.

Okuyama, K., M. Takayasu and H. Takayasu (1999), 'Zipf's Law in Income Distribution of Companies," *Physica A,* 269 (1), 125-131.

Park, K. and E. Domany (2001), "Power Law Distribution of Dividends in Horse Races," *Europhysics Letters,* 53 (4), 419-425.

Persky, Joseph (1992), "Retrospectives: Pareto's Law," *Journal of Economic Perspectives*, 6 (2), 181-192.

Robinson, William T. and Clauss Fornell (1985), "Sources of Market Pioneer Advantages in Consumer Goods Industries," *Journal of Marketing Research*, 22 (3), 305-317.

Simon, Herbert A. (1968), "On Judging the Plausibility of Theories," in B. van Rootselaar and J. F. Staal (Eds.), *Logic, Methodology and Philosophy of Sciences III*, Amsterdam: North Holland. 439-459.

Simon, Herbert A. (1955), "On a Class of Skew Distributions Functions," *Biometrika,* 42, 425-440.

Simon, Herbert A. and Charles P. Bonini (1958), "The Size Distribution of Business Firms," *American Economic Review,* 48 (4), 607-617.

Stanley, M. et al. (1995), "Zipf Plots and the Size Distribution of Firms," *Economic Letters,* 49 (4), 453-457.

Uncles, Mark, Andrew Ehrenberg and Kathy Hammond (1995), "Patterns of Buyer Behavior: Regularities, Models, and Extensions," *Marketing Science,* 14 (3), Part 2, G71-G78.

Urban, Glen L., Theresa Carter, Steven Gaskin and Zofia Mucha (1986), "Market Share Rewards to Pioneering Brands: An Empirical Analysis and Strategic Implications," *Management Science*, 32 (?), 645-659.

Zipf, George K. (1950), "A Note on Brand Names and Related Economic Phenomena," *Econometrica*, 18 (3), 260-63.

Zipf, George K. (1949), *Human Behavior and the Principle of Least Effort*, New York NY: Addison Wesley.