The Pareto Effect (80:20 rule) in Consumption of Beer, Wine and Spirits: A Preliminary Discussion

Keywords: Wine Marketing, Consumer Concentration, Pareto Share, Market Modelling

Cullen Habel
Wine Marketing Research Group
School of Marketing, University of South Australia
North Terrace, Adelaide SA 5000, Australia
Telephone: 61 (0)8 8302 0841, fax: 61 (0)8 8302 0442
cullen.habel@unisa.edu.au

Cam Rungie
Wine Marketing Research Group
School of Marketing, University of South Australia
North Terrace, Adelaide SA 5000, Australia.
Telephone: 61 (0)8 8302 0768, fax: 61 (0)8 8302 0442
cam.rungie@unisa.edu.au

Professor Larry Lockshin
Wine Marketing Research Group
School of Marketing, University of South Australia
North Terrace, Adelaide SA 5000, Australia
Telephone: 61 (0)8 8302 0261, fax: 61 (0)8 8302 0442
larry.lockshin@unisa.edu.au

Tony Spawton
Wine Marketing Research Group
School of Marketing, University of South Australia
North Terrace, Adelaide SA 5000, Australia
Telephone: 61 (0)8 8302 0714, fax: 61 (0)8 8302 0442
tony.spawton@unisa.edu.au

The Pareto Effect (80:20) in Consumption of Beer, Wine and Spirits: A Preliminary Discussion

Abstract

This paper considers two performance issues for several types of alcohol – category penetration and consumer concentration. Consumer concentration is addressed using the performance measure of "Pareto Share", which is defined as the percentage of category sales to the top 20% of its consumers. The beverage categories of beer, wine and spirits are first compared for their observed 1-week time period. The categories are then modelled, using the Negative Binomial Distribution in order to extrapolate market behaviour to longer time periods of observation – in this case a month and a year. Findings of this study are that the Pareto effect varies considerably across alcohol types and that the apparent Pareto effect increases as the sample time increases. The implications for managers are discussed and areas of further research highlighted.

Introduction

Categories have heavy and light users. In the first empirical study of this phenomenon, Twedt examined four categories of "Chicago Tribune" panel data in 1968, and asked the question "How important to marketing strategy is the 'heavy user'?" He found the "heavy half" accounted for more than 80% of the category purchases for the two categories he considered. A question of growing importance to alcohol beverage marketers must surely be "How important are *our* heavy users?" We may sagely nod at the mention of a "Rule of 80:20" where the top 20% of customers account for 80% of our sales volume but do we really know anything about it? Nagging questions include:

- Is there *80:20 rule* for wine and other alcohol beverage types, or is it simply a "rule of thumb"?
- What percentage of a category's consumption is due to the consumption of the heavy user? (How does consumption data present this?)

Other questions also come to mind once we begin analysing Pareto type distributions:

- Does the "rule" apply to each alcohol beverage category in the same manner?
- Does the time period of sampling matter? Is the Pareto Effect for a week's sample of consumer behaviour the same as a month's or a year's sample?

A search of the literature shows a few articles that have quantitatively addressed this question of buyer concentration ((Twedt 1968),(Schmittlein, Cooper and Morrison 1993), (Anschuetz 1997), (Rungie, Laurent and Habel 2002)) and none at all within the context of wine. Stanford (2000) has shown that wine consumers consume multiple alcohols but does not differentiate the heavy/ light consumption concentrations.

In the Australian context, wine consumption has evolved into a mass market orientation in the 1970's. This was due to an increase in its acceptance among women, a development in packaging technology (wine casks) which allowed bulk sales to less discriminating consumers and, among other things, a change in focus from the mature market of fortified wines to the growth markets of sparkling and table wines (Spawton 1991). Australia is predominantly a beer drinking country, with beer accounting for over 50% of alcohol consumption (Spawton 1995). Beer also accounts for the greatest percentage of men's alcohol consumption and wine accounts for the greatest percentage of women's consumption although the disparity is lower for wine (Spawton 1995). These results are consistent with the finding of this paper that wine has a greater penetration within the Australian community than beer.

The nature of this paper is not to describe the characteristics of heavy users. Even in the nascent field of wine consumption behaviour significant research has been conducted into that issue. Examples include (Keown and Casey 1995), (Blaylock and Blisard 1993), (Goldsmith and d'Hauteville 1998). These articles have all considered heavy users but in the sense of how heavy usage correlates to demographic, geographic and, interestingly, psychographic variables. This valuable in itself, but it differs from our research.

Our objective is more general. Many practitioners only have anecdotal evidence as to how many heavy users their categories have. Spirits marketers may think that a large amount of their volume is consumed by a relative few of their customers but have they ever demonstrated this? Wine marketers may like to think that their product has a broad appeal, so heavier users may account for less of the total consumption of wine, but how do they know?

Firstly, we aim remove some of the "hearsay" around this light/heavy user argument. Using real consumption data and modern modelling techniques, we answer the "what" question more than the "why". *What is* the Pareto Effect for Wine as compared to that of Beer or Spirits? Should it be found that the Pareto Share for wine is different to other categories it establishes that Pareto Share is not an absolute – that it varies between categories. We acknowledge the existing concept of the "Pareto Effect" and define "Pareto Share" as the percentage of sales volume that is made to the top 20% of a category's customers.

Secondly we aim to demonstrate that considering the customer base over a longer time period gives a dramatically different picture of customer concentration. Schmittlein, Cooper and Morrison (1993) first showed how buyer concentration (what we are calling "Pareto Share") increases as we consider longer time periods. We aim to demonstrate this finding within the context of alcoholic beverages and compare the nature of that relationship for three alcohol categories.

The Negative Binomial Distribution (NBD) has been shown to often represent the category repeat purchase rate across a population of shoppers. (Ehrenberg 1959). The fit is very strong when there is a high degree of stability in the purchase behaviour of the sample. Considering the maturity of the categories we are analysing we have no reason to believe that the NBD will be less applicable than it has been to soft drinks, aviation fuels, supermarket choices (Allsopp 2002), Detergents, Chocolate Bars, Biscuits, or Shampoos (Rungie, Laurent & Habel 2002) or any of the hundreds of categories the model has been applied to since 1959. Thus the NBD will be the major modelling tool we use to consider the "what ifs" of the category consumer profile. This ability to consider "what ifs" – especially with respect to time periods of observation is an established property of the Negative Binomial Distribution (Goodhardt and Ehrenberg 1967) (Schmittlein, Cooper and Morrison 1993).

The paper evolves as follows: i/ We consider the reported weekly consumption data for Wine, Beer and Spirits for a population of 4800 survey respondents, fit the NBD and assess the fit of the model to the observed data. ii/ We compare the Penetration and Pareto Share across product categories. iii/ The time period of observation is extrapolated within the model to simulate the change in perspective of a senior manager from one week to one month or one year. iv/ Theoretical and managerial implications are discussed.

Pareto Share

A corollary to the unevenness in concentration that comes about from the presence of heavy and light users is the commonly quoted "Pareto Effect". Pareto was a late 19^{th} century Italian engineer turned economist that first developed a mathematical description for inequity in a country's income distribution. (Persky 1992) His original empirical generalisation was later applied to areas such as statistical quality control and later to the social phenomenon of unevenness in customer concentration. (Weiner 2000)

To a marketing manager, Pareto means "80% of our sales are made to our top 20% of customers" (Buchanan 2002); (Sanders 1987). In reality, the proportion of sales to the top 20% of customers often seems closer to 60% and varies considerably, based on the time period of observation (Schmittlein, Cooper and Morrison 1993) and the market share of the brand (Rungie, Laurent and Habel 2002).

This percentage of sales to the top 20% of customers, or "Pareto Share" as we define it, appears to be a valuable tool to address the nature of product categories. It allows us to understand to what degree the heavy users of that category account for its turnover. "Pareto Share" may offer insight to brand managers as to whether to pursue increases in penetration of their brand as per (Ehrenberg and Goodhardt 1990), increases in loyalty (Zeithaml, Rust and Lemon 2001) or purchase frequency (Peppers, Rogers and et al. 1999) (Baldinger and Rubinson 1997).

Most importantly, looking at Pareto Shares gives us the opportunity to *compare* categories. Whilst there are a number of measures of behavioural loyalty, Pareto Share is one which is closely correlated to the shape parameter (the K parameter) of the distribution. Categories with similar Pareto Shares are dealing with customer bases that are very similar in their *behavioural* loyalty to the category. While penetration indicates the proportion of shoppers who have purchased the category one or more times, Pareto Share can add depth to an analysis by giving an indication of the *concentration* of these customers. Categories that may appear on the surface to be entirely dissimilar may have their customer bases behaving in the same sort of manner. Some early findings (Allsopp, 2002) show that - for example – "purchase of desserts" and "trips to the supermarket" are identical in how much volume is accounted for by the top 20% of their customers. This non-intuitive finding demonstrates the benefit of stepping back from the complexities and nuances of "what you see on the shelf" and using the tool of Pareto Share to clinically assess what your customer base is doing.

Results

Data Set

Survey data of approximately 4800 respondents was analysed. This was based on a quota sampling of the Australian population for the Australian Bureau of Statistics Population Monitor in 1995. Face to face interviews were taken in the house of the respondent. Whilst descriptive statistics for respondents were collected, it is the consumption statistics that will be analysed here. Respondents were asked to recall their consumption of alcohol products on a daily basis for the week prior to the interview.

Data has been recoded into weekly consumption (ml) of three alcohol types: Full Strength Beer, Wine (Comprising Red, White and Sparkling), and Spirits. A standardised drinks measure was then applied to convert the ml figure to a number of drinks, which is the data format required by the discrete NBD. The original data collection included a number of other liquor types that are not included in this analysis, namely liqueurs, fortified wines, light beers, extra light beers.

Observed Figures and Fit of the Model

We fitted the Negative Binomial Distribution to each of the categories. The NBD has a shape that will vary according to the two parameter values that are used in it. The act of fitting the distribution is a matter of determining which parameters create the shape of the NBD that most closely fits the observed data. The histograms of observed and theoretical consumption for each category were plotted on the same set of axes in order to a/ give a general picture of the category consumption behaviour and b/ indicate the fit of the model.

Figure 1: Wine has a high number of light users and the NBD fits well

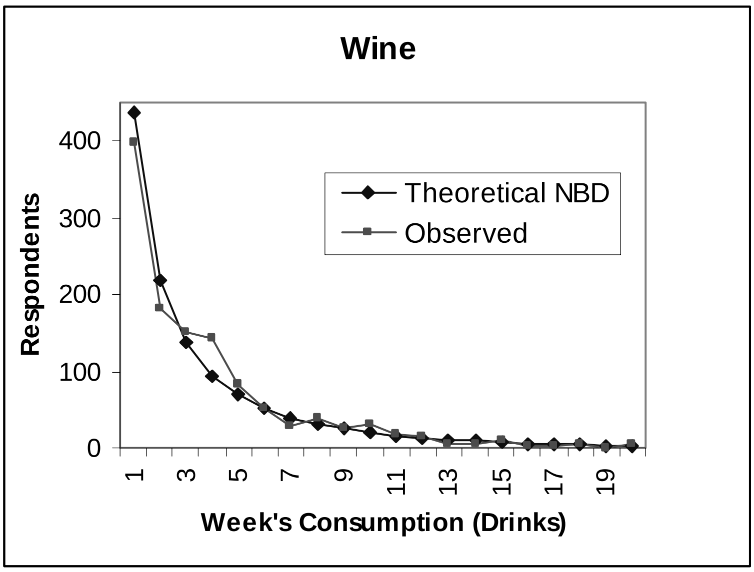

The first thing we see is that the theoretical figures fit the observed data quite closely. We also see that approximately 400 people consumed only one glass of wine in the week and that 200 people had only two glasses.

Figure 2: Beer has fewer light users and a poorer fit of the model.

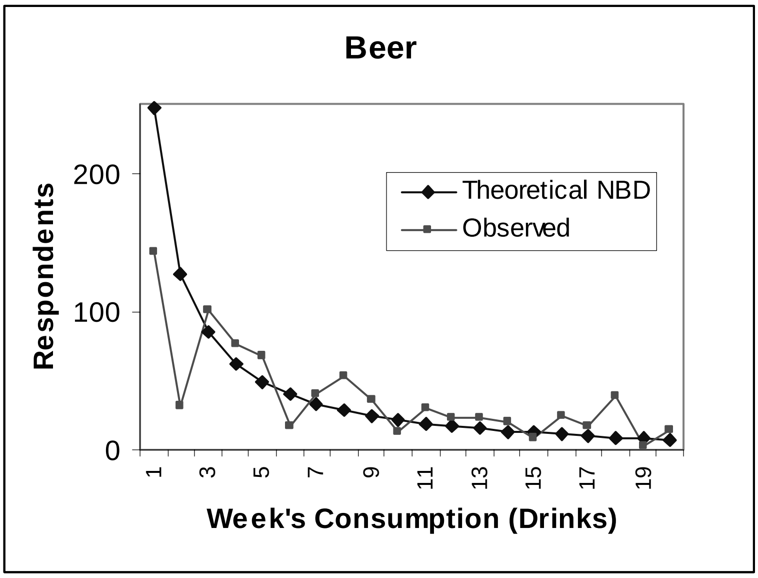

Aside from the poor fit of the model, it is valuable to note that the numbers of light users of the category is less than that for wine and that the “tail” to the right hand side of the graph is fatter. We will later show that this higher number of heavier users gives a higher Pareto Share.

Figure 3: Spirits also has fewer light users and a good fit of the model.

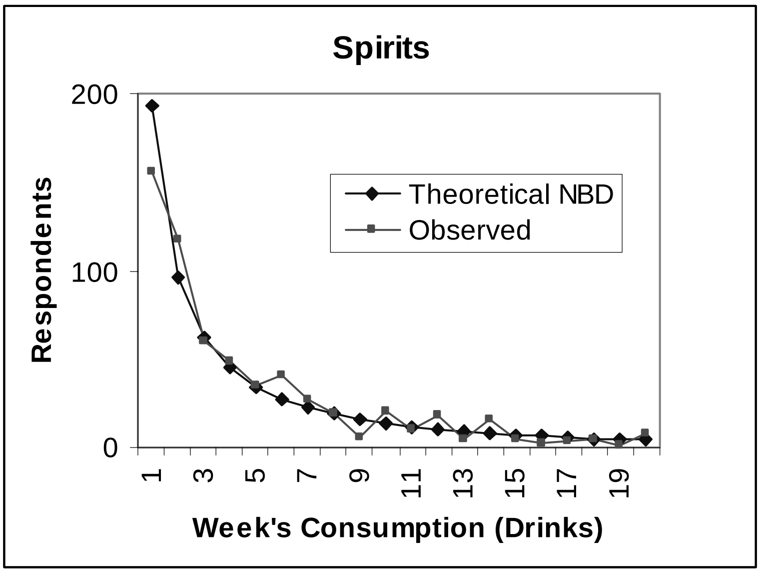

Spirits, again, exhibits the classic reverse J-curve of the NBD for certain parameter values. The NBD draws a line straight through the histogram and describes the category consumption quite well.

It is important to consider the beer category, in particular, with care. The fit is not as firm as the other two categories. The model overpredicts the two lightest consumption amounts (one per week and two per week). If, indeed, light beer was consumed more lightly than full strength beer, the observed numbers would be much closer to the theoretical line and the fit of the model would be much improved. Whilst the fit is poor at other points of the distribution, the worst "misfit" comes in the lightest consumption amounts. Therefore the poor fit for beer may be because we chose not to include light beer as part of the category or simply because of the limitations of the model in being able to fit all types of consumption and purchase behaviour. As this data is *reported* consumption the balance of irregularities *may* be due to respondents arriving at simple decision rules when answering the survey. For instance, one of the observed deviations occurs at a straightforward "one beer a day" decision rule.

Whilst looking at histograms of category consumption is valuable for considering the fit of the model and gaining a "picture" of the category, we need to look a little closer to gain insight as to the behaviour of the customer base.

The observed and theoretical Category Penetrations and Pareto Shares for the reported week are tabulated below. The Category Penetration is the proportion of the sample that has reported to have consumed the category one or more times in the reported period. Pareto Share is the percentage of the total category consumption accounted for by the top 20^{th} percentile of the customer base. Theoretical values are the same figures, calculated from the theoretical (NBD) estimates that we have graphed above.

Table 1: Wine, Beer and Spirits differ greatly in their observed penetrations and Pareto Shares

	Wine		Beer		Spirits	
	(O)	(T)	(O)	(T)	(O)	(T)
Penetration (%)	26	25	20	20	13	13
Pareto Share (%)	54	61	57	68	62	67

On considering the table, we see that 26% of the sample consumed wine in the sampled week. This was significantly above beer (20%) and spirits (13%). This may infer that wine has a more universal appeal than beer or spirits in Australia.

The observed *Pareto Share* for wine was also below that of both beer and spirits. That is, the top 20% of wine consumers accounted for proportionally less of the total consumption than in the case of beer or spirits. Therefore the results indicate that as well as having a broader appeal across respondents, wine appears to be drawing its consumption less from heavy users than the other two forms of liquor.

In all three cases the NBD overestimated the Pareto Share of the category. This could well be due to the exclusion of the more marginal product types for each category (fortifieds for wine, light beer for beer, liqueurs for spirits). A predominance of light users of these marginal forms would explain this consistent overestimation.

Differing Perspectives

Another factor clouding the Pareto Effect is likely to be the perspective of the observer. A sales manager with a twenty-year history is unlikely to view his customer base from the same perspective as the sales representative who started a month ago. The sales manager is likely to recall customers who simply have not purchased recently. These “once a year” customers are more likely to enter the sales managers’ mental sampling time and change his perspective in two ways: i/ he will see higher penetration (the people who consume at all in the period), and ii/ the heavy users (the “once a week” buyers) become a smaller percentage of the total customer base.

This intuitive phenomenon can be modelled. The longer mental period that a sales manager observes can be replicated by extrapolating the parameters of the model. Thus while we have observed a week of market behaviour and captured it in the model, we can now look beyond the observed data. If the sales manager tends to consider a twelve-month period, we can do the same thing with the help of the model.

Extrapolation

Schmittlein, Cooper and Morrison (1993) demonstrated that the NBD could be extrapolated to estimate the market behaviour for periods beyond the observed period of consumption. We decided to assess the Category Penetration and Pareto share for our three categories based on a consumption period of one month, and one year.

Category penetration over longer time periods

Firstly the category penetrations (theoreticals only) for each of the time periods are tabulated below.

Table 2: As the time period of observation increases the Penetrations increase.

Penetration	Wine (%)	Beer (%)	Spirits (%)
1 week	25	20	13
1 month	38	27	19
1 year	56	38	28

This straightforward result confirms our expectation that that sales manager will perceive a greater category penetration. It shows how the “once a year” customers will appear in the mental time frame (one year) of the sales manager and increase the overall penetration.

Whilst straightforward, the result is valuable in itself, in that it allows us to examine the penetration of wine on a long term basis. That is, in a year, close to 60% of people will drink *some* wine, compared to around 40% for beer and 30% for spirits.

Pareto Share Over Longer Time Periods

The Pareto Shares (theoreticals only) for the three categories, over the three time periods are plotted below.

Figure 5: As the time period of Observation Increases, the Pareto Share for the Category also Increases.

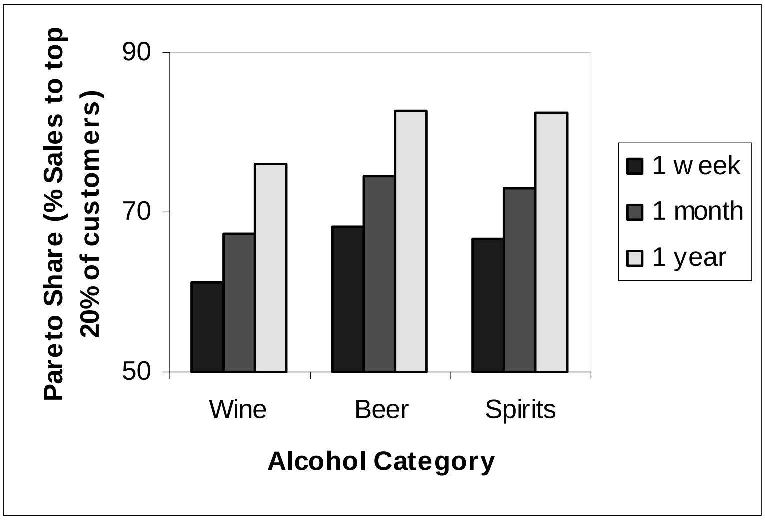

In all three cases the Pareto Share increases as the time period of observation increases. That is, if we consider wine consumption for one year, the top 20% of consumers would account for 75% of the wine consumed. Interestingly, both Beer and Spirits' Pareto Shares were about 80% for the one-year period which for these categories and that particular time period fits the terminology of "80:20 Rule". The variability, however, across both categories and observation periods indicates this is more of a "Rule of Thumb" than a common occurrence.

Discussion

What is Pareto and why does it vary?

It seems that the "Pareto Effect" is not a clear 80:20 relationship but that it varies among categories. Whilst it is not safe to make any general rule, the Pareto Share is more likely to be around 55% to 65% and most importantly – *it varies with the time period of observation.*

This variation in Pareto Effect can be described using a simple example. Consider a woman aged 23, her mother and her grandmother. The 23 year old woman consumes wine about twice a week as part of her busy social life, the mother maybe once a month at family get – togethers, and the grandmother maybe once a year at special occasions.

For a sampling time of a week, the 23-year-old gets into the sample but neither of the lighter users do. The sample is then full of consumers with a purchase frequency similar to hers. When we observe for a month, the mother is captured in the sample. By that time the daughter has consumed 8 times compared to the mother's one. Thus the concentration, or

Pareto Share of the one month sample is greater than the one week. Once the sampling period is extended to a year, the Daughter, Mother and Grandmother are all in the sample, with volumes of 104, 12 and 1 respectively. By including more light users the heavier users become a smaller percentage of the customer base and the Pareto share increases.

Is wine different?

The higher penetration of wine in Australia as demonstrated in Tables 1 and 2 may be due to the multiplicity of functions that wine performs, such as a vehicle for learning, play, cultural assimilation, socialising (Groves, Charters and Reynolds 2000). The broadness of the appeal of wine across sexes (Keown and Casey 1995) is likely to be a contributing factor at least in Australia.

At the longer time periods of observation the category penetration rates all increase, and wine's penetration increases its lead as we consider longer time periods. Wine consumption behaviour appears quite different to that of other forms of alcohol.

The lower observed Pareto Share for wine lends support to the notion that wine is the alcohol form that is consumed in moderation. The result may support the positioning of wine as an aid to health when consumed in moderation. (Spawton and Lockshin 2002)

One less reason to panic

The increase of Pareto Share over longer time periods of observation is predictable, and theoretically well substantiated. Managers need to note this in their assessment of their customer concentrations. There is a case of a senior brand manager becoming alarmed when they considered their customer concentration, to find later that they had been considering 3 year data and comparing it to one year data. Whilst it is rare that a sensible sales manager would explicitly compare 3 year data to one year data there may be an implicit assumption that Pareto Share is independent of the consumption period in question and this may lead to confusion when considering market structure.

A measure of the category's health?

Pareto Share is a measure of how beholden a category (in this case three alcohol beverage categories) is to their heavier users. If the wine category is seen as competing with other categories for a share of the market's consumption, it may well be argued that other forms of choice modelling techniques may be used to describe the choice behaviour. That is, perhaps Beer, Wine and Spirits may act as brands within the broader category of alcoholic beverages. In this case, when a consumer decides they are going to drink, they make the choice between a beer, a glass of wine, or spirits. A model for examining this choice behaviour is the Dirichlet (Goodhardt, Ehrenberg and Chatfield 1984), which describes the multivariate purchase rate for brands competing within a category and involves a significant amount of the NBD theory contained in this paper.

Under the Dirichlet there are a number of parameters, but only one that relates directly to the strength of that brand. Each brand what is termed its "Brand Alpha" parameter, which in the model is the "weapon" it takes into battle with the other brands. Increases in a customer's propensity to buy the brand in preference to others is directly represented as an increase in this "brand alpha". It has been shown that an increase in the brand alpha is positively associated with an increase in Pareto Share of the brand. (Rungie, Laurent and Habel 2002) Indeed, the shape parameter (K) that so closely correlates to Pareto Share in this paper would be replaced with a very similar "Brand Alpha" parameter if we were to extend the analysis to the choice between alcohol types under Dirichlet.

In this light, as a result of its direct relationship to brand strength, Pareto Share may well be the "signpost to the wine category's health" that we are seeking. A lower Pareto Share may indicate a stronger category.

Limitations and further research

We acknowledge the limitations of using face-to-face interviews, and relying upon respondent memory; this may have introduced a degree of collection bias. A cross sectional data set also has its limitations, as does the collection period of a single week which did not allow for observed/theoretical comparisons over differing periods of observation.

This paper has taken a single time period of observation and used the model to extrapolate to longer periods. There is no empirical verification of this extrapolation. Further research could include observed measures for other time periods of observation. As noted previously the exclusion of fortified wines, liqueurs and light beers may well have reduced the tightness of the fit of the NBD to our three categories.

This paper dealt with *consumption behaviour* – another consumer characteristic is *purchase behaviour.* A similar analysis conducted on repeat purchase panel data is likely to yield interesting results.

At brand level there are Pareto patterns that vary with market share as demonstrated in FMCG. (Rungie, Laurent and Habel 2002) The variation of Pareto Share for *wine brands* could be modeled using similar techniques. Considering the crowding of the wine category this would constitute an analysis at the lower ends of the market share spectrum. We find no evidence of this having been done before.

Conclusion

We have demonstrated how customer concentration varies between three categories of alcoholic beverage and how that concentration varies with the time horizon. As such this paper constitutes a replication of Scmittlein Cooper and Morrison's 1993 work in a different context. The result of different Pareto Shares across beer, wine and spirits demonstrates that Pareto Share can vary from category to category. We have introduced the terminology of "Pareto Share" and defined it as the percentage of sales made to the top 20% of customers, and demonstrated how it may serve as a guide to a category (or brand's) health. A deeper understanding of Pareto is likely to give the industry insight as to how to grow the category, and maybe give brand managers less cause to panic when they look at their panel data.

References

Allsop, J. (2002). Simple Quantitative Expression of Double Jeopardy. Marketing Science Centre. Adelaide, University of South Australia.

Anschuetz, N. (1997). "Profiting from the `80-20 rule of thumb'." Journal of Advertising Research **37**(6): 51.

Baldinger, A. L. and Rubinson, J. (1997). "The jeopardy in double jeopardy." Journal of Advertising Research **37**(6): 37.

Blaylock, J. R. and Blisard, W. N. (1993). "Wine Consumption by US Men." Applied Economics **24**: 645-651.

Buchanan, M. (2002). "Wealth Happens." Harvard Business Review **80**(4): 3.

Ehrenberg, A. S. C. (1959). "The Pattern of Consumer Purchases." Applied Statistics **8**(No. 1): 26-41.

Ehrenberg, A. S. C. and Goodhardt, G. J. (1990). "Double Jeopardy Revisited." Journal of Marketing **54**(3): 82.

Goldsmith, R. E. and d'Hauteville, F. (1998). "Heavy Wine Consumption: Empirical and Theoretical Perspectives." British Food Journal **100**(4): 184-190.

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

Groves, R., Charters, S. and Reynolds, C. (2000). "Imbibing, Inscribing, Integrating and Imparting: A Taxonomy of Wine Consumption Practices." Journal of Wine Research **Dec2000, Vol. 11 Issue 3**.

Keown, C. and Casey, M. (1995). "Purchasing Behaviour in the Northern Ireland Wine Market." British Food Journal **97**(1): 17-20.

Peppers, D., Rogers, M. and et al. (1999). "Is Your Company Ready for One-to-One Marketing?" **77**(1): 151.

Persky, J. (1992). "Pareto's law." Journal of Economic Perspectives **6**(2): 181.

Rungie, C. M., Laurent, G. and Habel, C. A. (2002). A New Model for the Pareto Effect (80:20) at Brand Level. Anzmac, Melbourne, Anzmac.

Sanders, R. (1987). "The Pareto Principle: Its Use and Abuse." Journal of Consumer Marketing **4**(1): 47.

Schmittlein, D. C., Cooper, L. G. and Morrison, D. G. (1993). "Truth in Concentration in the Land of (80/20) Laws." Marketing Science **12**(2): 167.

Soft Drinks, I. (2003). Soft Drinks Do What We Preach. Soft Drinks International.

Spawton, A. L. and Lockshin, L. (2002). Consumer Response to Societal Controls Designed to Modify Alcohol Usage and Abuse: Marketing Policy Implications. Bacchus to the Future Wine Conference.

Twedt, D. W. (1968). "How Important is Marketing Strategy to the 'Heavy User''?" Journal of Marketing **28**: 71-72.

Weiner, J. (2000). "Writing Off 80% of Customers Is Economics Gone Berserk." American Banker **165**(186): 11.

Zeithaml, V. A., Rust, R. T. and Lemon, K. N. (2001). "The Customer Pyramid: CREATING AND SERVING PROFITABLE CUSTOMERS." California Management Review **43**(4): 118.