

Using the Thiel Index Coefficient to Analyse Variation in Wine Consumption Habits

By Damien Wilson and Larry Lockshin

Both Authors are from
The Wine Marketing Research Group
School of marketing
University of South Australia
27-29 North Terrace
Adelaide SA 5000

Abstract

This study compares the reliability of using mean values to represent consumers' reported wine consumption behaviours with the use of a heterogeneity measure, the Thiel Index. We use the reported consumption of fortified wine from an Australian Population Monitor survey to segment consumers into age groups and then gender within age groups. The Thiel index is applied consecutively to the whole sample and then to each age group and gender in order to identify sub-segments, which are homogeneous enough to use the mean to represent the group. Because the Thiel Index has additive decomposability, the heterogeneity of the total sample can be broken down to find the sub-segments that are most homogeneous.

This property allows the market researcher to identify segments that have similar behaviour; or to locate sub-groups that have low consumption and, therefore, have more room to grow. It was found that mean consumption values are only applicable to most consumers within the 18-24 year old segment of fortified wine consumers. For 55+ year-old consumers of fortified wine, the level of heterogeneity is higher than any other age group. Even when

decomposed into the contributions by gender, both male and female consumers illustrate high degrees of heterogeneity. Therefore, use of mean values of consumption for these consumers would be inapplicable to many within the category of fortified wine consumption.

Introduction

Wine marketing is a relatively new discipline (Lockshin, 1997; Spawton and Bourqui, 1997), however, in recent years it has become topical. The literature on wine consumption is not well established because consumption data has been unavailable until recently. Some authors (James, 1978; Halliday, 1985; Beeston, 1994; Hardy and Roden, 1995) used production, apparent consumption, or sales data as a proxy for consumption data. These authors' interpretations on wine consumption have therefore been an inaccurate reflection of the actual wine consumption habits of the population. This practice has resulted in consumption summaries, but little is known about the variation in wine consumption habits of the population.

In 1998, the Australian Bureau of Statistics (ABS) captured data on a representative population sample of more than 6,000 individuals at the request of the Australian Wine and Brandy Corporation (Stanford, 1999, 2000). Stanford published summaries of the mean consumption habits by age groups and genders within this sample. The findings are the first real attempt to summarise the wine consumption habits of segments of the Australian population, based on reported consumption data.

Marketing practitioners use this information to make decisions on potential opportunities in the market. The problem is that the use of inaccurate or inapplicable market information is tantamount to guessing the opportunities in the market. The use of sales data was historically summarised into mean consumption values,

which have been used by marketers as proxies for the actual wine consumption habits of the population. In an effort to ascertain the applicability of using mean values for all members within age groups and by gender, the level of variation in responses, or heterogeneity as termed in econometrics, was measured with a tool called the Theil index coefficient.

The Theil index coefficient is adapted from the Theil index, a member of the family of entropy measures (Rungie and Juniper, 2001). The Theil index coefficient is the Theil index divided by its maximum value for the sample under analysis. The coefficient is used to return a number between one and zero, to indicate the absolute level of heterogeneity contained in any sample. Additionally, the tool is additively decomposable, meaning that the components within the sample can be broken into their contributory segments to overall heterogeneity. This advantage allows researchers to investigate the contributions to heterogeneity by sub-segments within a larger sample.

The layout of this paper is as follows: In the following section, we detail the wine consumption habits in Australia, including history, and the current knowledge. The next section details the problems of using traditional analysis methods to summarise wine consumption, and the consequences that marketers need to be aware of when using summarised market information. We then describe the Theil index coefficient, which was used to measure the variation in wine consumption, and why the marketer, to resolve potential problems around using summarised information, should use this tool. The next section details the findings of our analysis. Following are the results from comparing the mean and the Theil index coefficient to analyse fortified wine consumption habits of a population representative sample of Australians. Finally we present the conclusions from our analysis, including recommendations, shortcomings and suggested further research.

Wine Consumption

This section summarises the key literature on wine consumption in Australia, the problems associated with investigating wine consumption, and the recent change to using reported wine consumption data, instead of using anecdotal or unreliable proxy measures.

Historically, records of wine consumption in Australia are based on anecdotal and sales data, including 'apparent consumption'. Beeston (1994) has written extensively on the history of wine in Australia (Beeston, 1994). Beeston describes the patterns of wine consumption in Australia, as being predominately fortified wine from European settlement until the late 1960s and early 1970s, when table wine production increased markedly. This significant change in production patterns was attributed to the adoption of refrigeration in wine making practice. Other authors, such as James Halliday, Hardy and Roden, have relied on anecdotal and apparent consumption to assert their opinions on wine consumption habits (James, 1978; Halliday, 1985; Hardy and Roden, 1995).

These authors assert that once table wine consumption began to increase, white wine consumption became the preferred choice for most wine drinkers (Beeston, 1994; Buckley, 1996; Trueland, 2000). This suggested growth in wine consumption has been linked to the growth in women consuming wine, although the assertions are not supported with data. Further, the mid-1990s have witnessed the growth in red wine consumption. Again, this finding is primarily supported by sales data. However, the ABS instigated studies in the mid to late 1990s, on reported consumption habits, which add support to this assertion.

It is only with the gathering of information on reported alcohol consumption of the population that Stanford and Spawton have started to investigate patterns of wine consumption (Stanford,

1999; Spawton, 2000; Stanford, 2000; Spawton, 2001). Stanford was one of the first authors to use consumption data to present descriptive statistics on consumption habits in Australia (Stanford, 1999, 2000). Stanford, in both papers, presented aggregate information on the population, and the mean values for age, gender and other segments, in an effort to convey valuable descriptive statistics on the consumption habits of the population.

The source of information in Stanford's studies was from the Australian Bureau of Statistics' population survey monitor (PSM). The PSM is a survey conducted periodically, in an effort to determine the habits of the Australian population, by sampling a representative number of consumers. Responses from the sample are weighted to reflect the respondent's representation in the Australian population based on the age, gender and location of the respondent. Consequently, the outcomes of the study are expected to closely reflect the actual habits of the entirety of the Australian population.

Issues in Forecasting Wine Demand

Traditionally, mean values have been used to describe patterns of wine consumption, but these have been based on apparent consumption, based on sales or production data (James, 1978; Halliday, 1985; Beeston, 1994; Hardy and Roden, 1995). Using apparent consumption, sales and/or production data to assert patterns in consumption habits is erroneous. Despite obvious correlations between the volume of wine produced and consumed, the use of production figures to assert consumption, does not consider the inventories held by wine retailers. Additionally, retail sales figures do not consider the consumer's predisposition to mature wine before consumption, gift giving, or even sharing the same wine with other consumers. Consequently, it is virtually impossible to accurately investigate patterns in consumption habits, without the existence of consumption data.

Prior to the existence of consumption data, the 'best guess', opinion on consumption habits, personal preference or production capacities were used to determine how much of each wine category was produced from one vintage to the next. Estimating wine demand was not a serious concern for wineries over the past 10 years, as demand was perennially in excess of supply and most wine would sell out before the next vintage had commenced (Awbc, 2002).

With 2002 being the first year of excess supply in about 20 years (AWBC, 2002), wineries are looking for ways to clear their pending excess wine production. Adding to wineries' concerns are that many growers have continued to plant grapevines at an alarming rate, up until last year. Many of these new and young vines are still to produce a commercial crop, and will continue to add to the industry's wine production for the next ten years (vintage conditions notwithstanding). Consequently, marketers are investigating ways to cater to consumer demand more effectively.

Data analysts of wine consumption data like Lawrie Stanford of the Australian Wine and Brandy Corporation (1999, 2000) are relied upon for providing useful wine consumption data. Marketers within the wine industry use descriptive statistics provided by the AWBC as the foundations for developing marketing strategies. Therefore, it is important that the information provided to the industry is of the most practicable value.

Stanford's findings on mean wine consumption for age groups and genders (1999, 2000) are an important tool for indicating the approximate demands on wine production from year to year. By extrapolating the mean consumption from Stanford's research into a firm's production categories, a producer aims to forecast approximate demand for their wine in the coming year(s).

Information Summaries

Malhotra, Hair and other authors (Malhotra, 1993; Hair, Anderson, Tatham and Black, 1995) support the premise that data can be summarised with descriptive statistics of central tendency, like the mean and standard deviation. Descriptive statistics are useful because they provide practical implications from an otherwise, unwieldy data set.

However, descriptive statistics of central tendency rely on a data set being normally or uniformly distributed. The level of variation in consumption indicates how consistent different summary measures are for applying to individuals and smaller segments within a larger set (Ehrenberg, 1975; Wannacott and Wannacott, 1990). Marketers use the normal distribution to consider the possible number of potential errors on the production line, or for forecasting the potential demand for a new product based on the findings of a small and representative sample of consumers.

However, an assumption of a normal distribution is questionable, in light of the literature on market patterns by authors such as Ehrenberg (1993, 1997, 2000), Wright *et al.* (1998) and others. These authors espouse that most consumer market distributions follow the pattern of a negative binomial distribution (NBD) (Ehrenberg and Bound, 1993; Ehrenberg, 1997; Wright, Sharp and Sharp, 1998; Sharp and Wright, 1999; Ehrenberg and Uncles, 2000; Romaniuk and Sharp, 2000). This distribution has demonstrated, across a variety of markets, that most product/service/consumption habits follow an NBD, which has a small number purchasing larger amounts and a long tail of light or irregular buyers of a brand or category. For a negative binomial distribution, the mean and standard deviation would not have the descriptive accuracy that is possible with a normally or uniformly distributed population.

Heckman noted that variation exists within a sample and should be measured. Specifically, he said that:

“Because the source of sample variation critically affects the formulation and estimation of many models of discrete choice.....the formulation of an economic model for a ‘typical individual’...is absent from many of the best papers in this literature” (Heckman in (Hensher and Johnson, 1981), i).

Consequently, the goal of this study is to use the Theil index coefficient to measure the variation within a sample, in an effort to illustrate the sample’s uniformity or lack thereof, in consumption. The findings will highlight any need to segment the market further, in order to identify relatively uniform segments of wine consumption, which can be more effectively catered to by wine producers.

Measuring Sample Variation

Thiel proposed an index, which measured heterogeneity in econometrics (Thiel, 1967). Thiel’s index was developed to quantify a data set’s heterogeneity in a descriptive statistic. The Theil index incorporates the higher moments of a data set in its calculations, effectively accounting for the heterogeneity in each subject’s responses (Conceição and Ferreira, 2000). The result is that the Theil index effectively summarises the heterogeneity within a data set, indicating an index of variation.

Thiel’s index number has primarily been used for examining and describing the changing patterns of income inequality over time (Galbraith, 1997; Conceição and Ferreira, 2000; Conceição, 2001), but has also been applied to the health industry (Castillo-Salgado, Schneider, Loyola, Mujica, Roca and Yerg, 2001). However, an examination of the literature suggests that the Thiel index has not

been used in marketing (Bearden, Netemeyer and Mobley, 1993; Bruner and Hensel, 1998).

The Theil index is a heterogeneity measure that provides a measure of the discrepancies between the distribution of a resource within a population (Theil, 1967; Sheret, 1991; Galbraith, 1997; Conceição and Ferreira, 2000). Essentially, “the Theil index presents a number that illustrates heterogeneity by comparing the distribution across populations through summing across groups, the weighted natural logarithm of the ratio between the share of the population for each group characteristic being measured” (Conceição and Ferreira, 2000, pg. 13). The equation is provided below.

Equation 1- Equation used to generate the Theil Index for any data set

$$T = \sum_{i=1}^n [p_i (\ln(p_i) - \ln(w_i))]$$

In the equation above, p represents the proportion of the total resource [in this case, wine] consumed by each respondent; and w indicates the proportion that each respondent represents in the sample. The Theil index is calculated by multiplying each group being analysed by the proportion of overall consumption. In this way, the proportion of wine consumed weights precisely each group.

By calculating the Theil index, the researcher is provided with an easily compared value against other sets within the sample being analysed. However, the shortcomings of the Theil index are that the sets under analysis must be mutually exclusive and completely exhaustive [eg. male/female, age groups, income brackets]. Secondly, the value returned for the Thiel index has no absolute scale, as it is relative to the set under analysis. The index returned for each set is dependent on the number of individuals within the set (Theil, 1967; Shorrocks, 1980; Galbraith, 1997; Galbraith and Jiaqing, 1997; Conceição and Ferreira, 2000).

Sheret (1991) discussed the potential to use the Theil index as an absolute value, by presenting it as a percentage of the maximum heterogeneity possible in a sample. The subsequent value is expressed as a coefficient of heterogeneity, which is bound between the values of zero and one. Consequently, comparisons across different sets of data can be made, despite the different number of individuals within each data set.

The aim of this paper is to communicate a method that measures the level of heterogeneity in a market, identifying the need for a market to be segmented into categories that illustrate relatively uniform demand for fortified wine.

Methodology

The data from Stanford's (1999, 2000) paper on alcohol consumption habits was used for this application of the Theil index. In order to illustrate the benefits of this tool, and in the interests of brevity, applying the Theil index to fortified wine consumption provides an illustration of the additive decomposability of the Theil index coefficient.

The formula below is an application of the Theil index formula to the characteristics of the sample provided from Stanford (1999, 2000), by the ABS. The equation illustrates the way that the Theil index differentiates the groups by multiplying each group being analysed by the proportion of overall consumption contributed by that group. In this way, the proportion of alcohol consumed weights precisely each group (Theil, 1967; Shorrocks, 1980; Conceição and Ferreira, 2000).

Equation 2: Equation used to generate the Theil Index for fortified wine consumption

$$T_{Total} = \sum_{i=1}^n [\%consumed_i (\ln(\%consumed_i)) - (\ln(\%population_i))]]$$

The result provides us with a Theil index value of heterogeneity for fortified wine consumption. The value of the Theil index is then divided by the upper bound for fortified wine consumption. The upper bound is calculated by taking the natural logarithm of the number of fortified wine consumers in the sample. The quotient is presented as a percentage of heterogeneity and is termed the Theil index coefficient.

The Theil index coefficient can then be decomposed into the contributions to heterogeneity by each age group. In order to calculate the total contribution to heterogeneity, the value within each group must be calculated and summed with the heterogeneity between each group as shown in formula 3 below (Galbraith, 1997; Conceição and Ferreira, 2000).

Equation 3: Formula for calculating the Theil index by decomposing into within and between group heterogeneity

$$T_{Total} = \sum_{i=1}^n T'_{within_i} + T'_{between}$$

The characteristic of additive decomposability (Rungie and Juniper, 2001) means that the weighted contributions for each age group, and the heterogeneity between each age group, will sum to the Theil index value for fortified wine in this example. The proof that the calculations in the following formulas can be found in (Galbraith, 1997; Conceição and Ferreira, 2000). However, for this study, only the pertinent formulae are provided.

The formula for calculating the heterogeneity within each age-group is illustrated below in Equation 3.

Equation 3: Formula for calculating contributions to the Theil Index for heterogeneity within age-groups

$$T_{Within} = \frac{population_i}{population_{fortified}} \sum_{i=1}^n [\%consumed_i (\ln(\%consumed_i)) - (\ln(1/population_i))]]$$

The contributions to the Theil index for heterogeneity within age-groups must be weighted by the age-group's 'volume consumed' representation in the larger group under analysis.

In order to calculate the Theil index for each decomposition, the between group heterogeneity value is calculated using the formula indicated below. weighted within age-group heterogeneity is summed to give the Theil index for the age group. The formula for generating the Theil index from sub-groups is given below in Equation 4.

Equation 4: Formula for calculating between group heterogeneity for the Theil Index

$$T_{Between} = \sum_{i=1}^n [\%consumed_i (\ln(\%consumed_i)) - (\ln(\%population_i))]]$$

In the equation above, the subscript 'i' is to indicate the totals for each age group as opposed to each case. In this way, the heterogeneity captured between each age group is calculated in addition to the heterogeneity within each category.

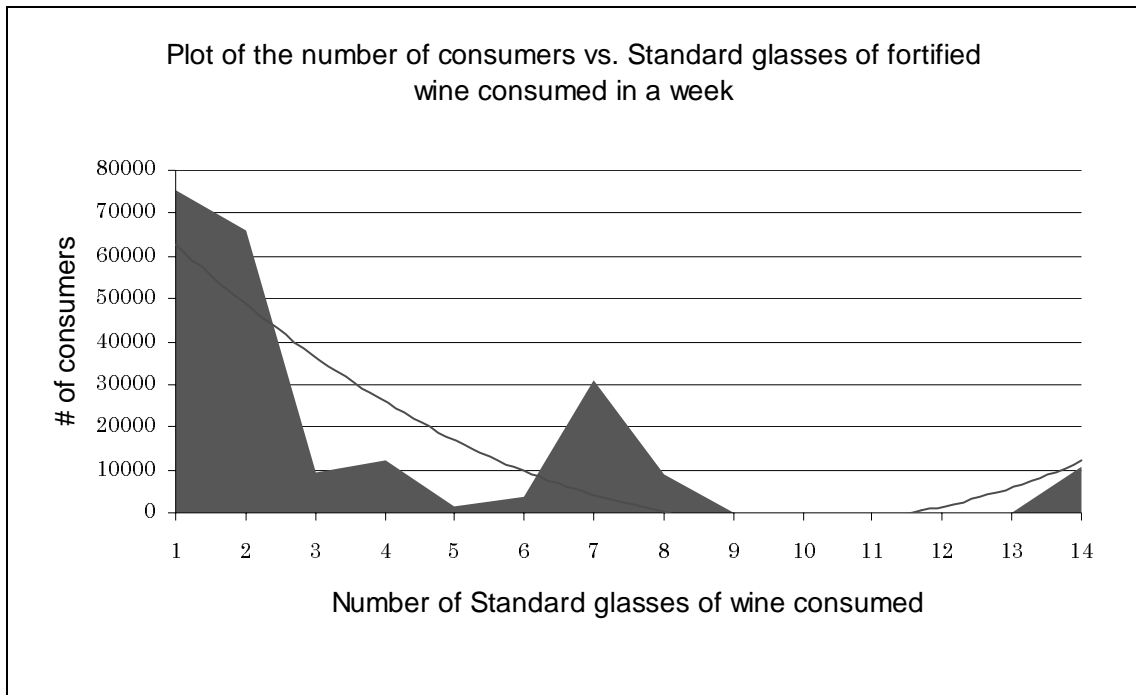
The process of decomposition is repeated for genders within each age group. Once Theil indices have been calculated for each gender's contribution to an age group's heterogeneity, the contribution to the Theil index coefficient for each age group is weighted by each gender's proportion of within group heterogeneity. This is an effective process for illustrating the contributions to an age group's heterogeneity by each gender. This action will identify which gender has most uniform consumption habits, and which ones contribute most to wine category heterogeneity. The higher the heterogeneity, the greater

the proportion of light users in the category, and the less descriptive the mean value is for individuals within the category.

Findings

Figure 1 below illustrates the alcohol consumption patterns in the form of a histogram. The 'x' axis indicates the number of standard drinks of fortified wine consumed in the week prior to the survey. The 'y' axis indicates the number of consumers who drank the corresponding volume of wine in the week prior to the study, according to the weightings provided by the ABS.

Figure 1 Plot of the number of consumers vs. the number of standard fortified wine drinks consumed in a week in 1998



The shape of the histogram is different from the expected bell curve of a normal distribution, or a straight line of a uniform distribution. With the trend-line applied to the histogram, it is impossible to make an argument for a normal or uniform distribution.

Table 1 Comparing coefficients of heterogeneity for wine consumption by category in Australia in 1998

	Weekly Mean drinks	Theil Index coefficient
Fortified wine consumers	3	68%

Table 1 shows that the mean consumption of fortified wine is 3 drinks per week for those consumers who drank fortified wine. However, this mean value only captures approximately 10,000 consumers, from a population of approximately 220,000 fortified wine drinkers, as shown in Figure 1. Consequently, the mean value has little relevance for applying to those within this wine category.

The Theil index coefficient shows that there is evidence of heterogeneity within this category, and that segmentation of this group should occur, to identify segments with more uniform consumption habits.

The Theil index coefficient of 68% for this wine category was deconstructed into the segments that contribute to this value. As Stanford's (1999, 2000) articles separated the groups of wine consumers into age group and gender, this study followed that pattern. In this way, those who wish to may compare the findings of Stanford with the outcomes of this study.

Table 2 introduces the concept of the weighted Theil index coefficient. After calculating the Theil index coefficient for each category under analysis, the researcher aims to identify the contribution to heterogeneity by sub-groups within each category. As such, each of these groups is mutually exclusive and completely exhaustive, as mentioned earlier in text. When investigating the contribution of each sub-group to a category's heterogeneity, it is imperative to weight each sub-group by the proportion of consumers in each sub-group compared to the total number of consumers within the category.

Table 2: Descriptive data from 1998 on fortified wine consumption by age groups.

Age groups	Weekly Mean drinks	Theil Index coefficients	Weighted Theil Index coefficients
18-24	1	82%	1%
25-34	4	87%	13%
35-44	2	76%	10%
45-54	4	83%	14%
55+	4	70%	31%

Table 2 shows that the mean number of standard drinks of fortified wine varies between one and four, increasing with age, except for the 25-34 year old consumers. The 25-34 year old consumers exhibit the most heterogeneity within their group, for all age groups, using the Theil index coefficient. However, in the context of all fortified wine consumers, the contribution from 18-24 year old consumers is minute. Their consumption habits are relatively uniform with respect to all others in the category.

Again, referring to Figure 1, the mean values do not reflect the consumption habits of most consumers within the category. However, the deconstruction into age groups shows that more information about the sample is provided. Further, the age group with the highest heterogeneity, 55+ year olds, will be deconstructed further, to show the contributions by gender within the age group.

This age group consumed 352,814 Litres of fortified wine in the week prior to this study. Approximately 92,434 consumers were responsible for this amount of fortified wine.

The Theil index coefficient for 55+-year-old fortified wine consumers was calculated as 70%, as shown in Table 2. When the weighting is applied for the number of 55+ year old consumers, 31 percentage points of the total 68% in the fortified wine drinking population, is the contribution to fortified wine heterogeneity. The contribution to the age group's heterogeneity from each gender is indicated below.

Table 3 Descriptive statistics of heterogeneity by gender within 55+ year old fortified wine consumers

55+ years	Weekly Mean drinks	Theil Index coefficients	Percentage point contribution to category
Males	4	29%	12.8%
Females	4	41%	17.9%

There is little difference in heterogeneity between genders. Therefore, the proportion of each gender in this age group is responsible for about the same proportion of fortified wine consumed. Additionally, the Theil index coefficients sum to give a total of 70%, which equals the calculation for this age group in Table 2.

Conclusions

Our analysis of the consumption habits for the category of fortified wine drinkers found that consumers under the age of 25 are found to have very uniform fortified wine consumption habits. Based on these findings, estimating the volume demands of 18-24 year old fortified wine drinkers would be easier than estimating any other age group of fortified wine drinkers. This is because this age group was identified as a segment that exhibits relatively uniform consumption habits.

This contrasts with the results for consumers over the age of 55, who were found to have very heterogeneous fortified wine consumption habits. This finding indicates that the over 55-age group should be segmented further, in an effort to identify groups with more uniform consumption habits. Perhaps there are a smaller group of older consumers who are abusing fortified wine and further marketing would have negative social implications. On the other hand, there is a larger group that drinks very little fortified wine, who might be marketed to.

Applications of this Study

The Theil index coefficient has applications for use in academic and practical investigations. For academics, the Theil index

coefficient can be used to analyse market data, and support a need for further segmentation, or support the use of mean values for the majority of consumers within the sample. The strength in the tool is that the Theil index coefficient determines whether a segment of the market has relatively uniform consumption habits, or is in need of further segmentation. Additionally, the tool can be used for any discipline where uniformity [or lack thereof] is a test of compliance or reliability.

The Theil index coefficient has been shown to have a practical application in wine marketing when other measures of variation are inapplicable, unworkable or inappropriate. Although the investigator is only able to measure the variation in mutually exclusive and completely exhaustive [MECE] groups within a data set. However, the fact that the application can be used in a field outside of its usual application is positive for other researchers.

Historically, the common practice for using the Theil index was to compare the same samples over time to illustrate change. The application in this thesis shows that there is potential to use the tool for its heterogeneity measurement as well as its temporal capacity. However, for disciplines where consistency over time is necessary to support theory, the Theil index coefficient can be used to illustrate how uniform a group's behaviour is over any period of time.

Additionally, the use of the Theil index coefficient in a wine marketing application provides an extension of its use for indicating financial inequality between and within countries. The Theil index coefficient highlights each group's heterogeneity in wine consumption. Further, the Theil index coefficient is used to illustrate the segments within each group, that are responsible for the bulk of each group's heterogeneity. The Theil index coefficient's value at illustrating the contributions of sub-groups to overall heterogeneity is where this tool has advantages over other measures of variance.

The value to managers is that the greater the heterogeneity in a sample's wine consumption, the more difficult it is to apply descriptive statistics of central tendency to those within the sample, as the sample is less likely to adhere to a normal or uniform distribution. Consequently, marketers who use the Theil index for analysing their own data will be made aware of the segments, and sub-segments, of their market that present opportunities for growth.

Practitioners can use the Theil index coefficient to determine whether mean values can be attributed to a majority of consumers in their market. Where high values exist for the Theil index coefficient, the marketer is made aware that there is substantial heterogeneity in the sample, and that a small number of consumers are responsible for a large volume of sales or consumption of a product or service. Consequently, a large heterogeneity value indicates that further segmentation of a segment should occur, if the marketer aims to have mean values that are an accurate reflection of the habits of the sample under analysis. Or there is an indication of a large number of light consumers, who may constitute an opportunity for growth.

This use of the Theil index is outside of the realm of income inequality where it has been used previously. However, it is clear from the results of the study that the Theil index has an application for indicating the need for further segmentation in a market. Additionally, the Theil index can be used to confirm the consistency of behaviours in a segment, and provide support for applying the mean value for a sub-segment to those within the segment.

Further Research

The use of this inequality index has been shown to have an application for wine consumption. However, this study should be replicated, to improve the tool's power. The original use of the Theil index was for illustrating changes in inequality over time. Both practitioners and academics can use this tool to illustrate how effectively the needs of a market have been met over time.

As the index is sensitive to changes over time, it would be a useful tool for monitoring changes in wine consumption patterns, particularly for segments within a population. Additionally, the segments with high heterogeneity contain the greater numbers of infrequent and low volume consumers. According to the findings of Sharp and Allsopp (2002) and Anschuetz (2002), a high number of light users are indicative of potential opportunities for growth (Anschuetz, 2002; Sharp and Allsopp, 2002). By monitoring the changes in indices over time, segments can be monitored for growth in consumption.

References

- Anschuetz, N. (2002). Why a brand's Most Valuable Customer is the Next One It Adds. *Journal of Advertising Research* 42(1): 3.
- AWBC, A. W. a. B. C. (2002). Global Wine- Wine in Perspective, Australian Wine and Brandy Corporation. 2002.
- Beeston, J. (1994). A Concise History of Australian Wine. Sydney, Australia, Allen & Unwin: 336.
- Buckley, K. (1996). Euromonitor Weighs the Near-term Market, Wine Business Monthly. 1997.
- Conceição, P. and P. Ferreira (2000). The Young Person's Guide to the Theil Index: Suggesting Intuitive Interpretations and Exploring Analytical Applications, UTIP. 2001.
- Ehrenberg, A. S. C. (1975). Data Reduction. Bristol, J.W Arrowsmith: 391.
- Ehrenberg, A. S. C. (1997). Description and Prediction. *Journal of Advertising Research*(November-December): 17-22.
- Ehrenberg, A. S. C. and J. A. Bound (1993). Predictability and Prediction. *Journal of the Royal Statistical Society Association* 156(Part 2): 167-206.
- Ehrenberg, A. S. C. and M. D. Uncles (2000). Understanding Dirichlet-Type Markets. London, South Bank University and University of New South Wales (www.sbu.ac.uk/RandDI/Publications).
- Galbraith, J. K. (1997). The Distribution of Income, LBJ School of Public Affairs. 2001.
- Galbraith, J. K. and L. Jiaqing (1997). Cluster and Discriminant Analysis on Time-Series as a Research Tool, LBJ School of Public Affairs. 2002.
- Hair, J. F. J., R. E. Anderson, R. L. Tatham and W. C. Black (1995). Multivariate Data Analysis. New Jersey, Prentice Hall: 131-133.
- Halliday, J. (1985). The Australian Wine Compendium. Sydney, Angus and Robertson: 576.

- Hardy, T. and M. Roden (1995). Australian Wine- A Pictorial Guide. Adelaide, Pictorial Wine Publishers: 352.
- Hensher, D. A. and L. W. Johnson (1981). Applied Discrete Choice Modelling. New York, Halsted Press: 468.
- James, W. (1978). Wine in Australia. Melbourne, Sun Publishing house: 156.
- Lockshin, L. (1997). Marketing and Management for the Wine Industry. *Australian and New Zealand Wine Industry Journal* **12**(1): 71-72.
- Malhotra, N. K. (1993). Marketing Research- An Applied Orientation. New Jersey, Prentice Hall: 857.
- Romaniuk, J. and B. Sharp (2000). Using Known Patterns in Image Data to Determine Brand Positioning. *International Journal of Market Research* **42**(Issue 2, Spring/Summer): 219-230.
- Rungie, C. and J. Juniper (2001). Entropy Paper. D. Wilson. Adelaide, University of South Australia.
- Sharp, B. and J. Allsopp (2002). Where Does Sales Revenue Growth come from? ANZMAC 2002, Melbourne, Deakin University- Burwood Campus.
- Sharp, B. and M. Wright (1999). There are Two Types of Repeat Purchase Markets. 28th European Marketing Academy Conference, Berlin, Germany, Institute of Marketing, Humboldt-University.
- Sheret, M. (1991). Measuring Inequality in Education. Conference of the Australian Association for Research in Education, Sydney.
- Shorrocks, A. F. (1980). The Class of Additively Decomposable Inequality Measures. *Econometrica* **48**: 613-625.
- Spawton, A. (2000). A Taxonomy of Wine Consumer Research in Australia 1990-2000. XXIVeme Congress of Wine and Vine, Paris, France, OIV.
- Spawton, A. (2001). Alcohol Consumption Habits. Division of Business and Enterprise. Adelaide, University of South Australia: 500.

- Spawton, A. and M. Bourqui (1997). Consumers Reactions to Risk Reducers. *Recherche* **97**(1): 3-17.
- Stanford, L. (1999). Observations on Alcohol Beverage Consumption. Wine Industry Outlook Conference, Adelaide, Australian Wine and Brandy Corporation.
- Stanford, L. (2000). Observations on Alcohol Beverage Consumption. *Australian and New Zealand Wine Industry Journal* **15**(1): 14-27.
- Theil, H. (1967). Economics and Information Theory. Chicago, Rand McNally and Co.
- Trueland, J. (2000). New social status sends women off to the bar. *The Scotsman*. Edinburgh: 7.
- Wannacott, T. H. and R. J. Wannacott (1990). Introductory Statistics for Business and Economics. Singapore, Wiley: 815.
- Wright, M., A. Sharp and B. Sharp (1998). Are Australasian Brands Different? *Journal of Product and Brand Management* **7**(6): 465-480.